在线学习资料支持

您可以在华为企业业务网站获得E-Learning课程、培训教材 、产品资料、软件工具、技术案例等：

1、 E-Learning课程：登录*华为在线学习网站*，进入 “*华为培训/在线学习*” 栏目

免费E-Learning课：对网站所有用户免费开放

职业认证E-Learning课：通过任何一项职业认证即可学习所有职业认证培训E-Learning课程

渠道赋能E-Learning课：对华为企业业务合作伙伴免费开放

2、培训教材：登录*华为在线学习网站*，进入 “*华为培训/面授培训*” ， 在具体课程页面即可下载教材。

华为职业认证培训教材、华为产品技术培训教材。无需注册即可下载

3、华为在线公开课(LVC)： http://support.huawei.com/ecommunity/bbs/10154479.html

企业网络、UC&C、安全、存储等诸多领域的职业认证课程，华为讲师公开授课

4、产品资料下载： http://support.huawei.com/enterprise/#tabname=productsupport

5、软件工具下载：http://support.huawei.com/enterprise/#tabname=softwaredownload

更多内容请访问：

- http://learning.huawei.com/cn
- http://support.huawei.com/enterprise/
- http://support.huawei.com/ecommunity/

华为认证系列教程

HCDP-IERN

部署企业级路由网络

实验指导书

华为技术有限公司

华为认证系列教程

HCDP-IERN部署企业级路由网络

实验指导书

第1.6版本

华为认证体系介绍

依托华为公司雄厚的技术实力和专业的培训体系，华为认证考虑到不同客户对ICT技术不同层次的需求，致力于为客户提供实战性、专业化的技术认证。

根据ICT技术的特点和客户不同层次的需求，华为认证为客户提供面向十三个方向的四级认证体系。

HCNA(HCDA)认证定位于中小型网络的基本配置和维护。HCNA(HCDA)认证包括但不限于：网络基础知识；流行网络的基本连接方法；基本的网络建造；基本的网络故障排除；华为路由交换设备的安装和调试。通过 HCNA(HCDA)认证，将证明您对中小型网络有初步的了解，了解面向中小型企业的网络通用技术，并具备协助设计中小企业网络以及使用华为路由交换设备实施设计的能力。拥有通过 HCNA(HCDA)认证的工程师，意味着中小企业有能力完成基本网络搭建，并将基本的语音、无线、云、安全和存储集成到网络之中，满足各种应用对网络的使用需求。

HCNP-Enterprise (HCDP-Enterprise)认证定位于中小型网络的构建和管理。HCNP-Enterprise (HCDP-Enterprise)认证包括但不限于：网络基础知识；交换机和路由器原理；TCP/IP 协议簇；路由协议；访问控制；网络故障的排除；华为路由交换设备的安装和调试。通过 HCNP-Enterprise (HCDP-Enterprise)认证，将证明您对中小型网络有全面深入的了解，掌握面向中小型企业的网络通用技术，并具备独立设计中小企业网络以及使用华为路由交换设备实施设计的能力。拥有通过 HCNP-Enterprise (HCDP-Enterprise)认证的工程师，意味着中小企业有能力完成完整网络的搭建，将企业中所需的语音、无线、云、安全和存储全面地集成到网络之中，并且能满足各种应用对网络的使用需求，进而提供较高的安全性、可用性和可靠性。

HCIE-Enterprise 认证定位于大中型复杂网络的构建、优化和管理。HCIE-Enterprise 认证包括但不限于：不同网络和各种路由器交换机之间的互联；复杂连接问题的解决；使用技术解决方案提高带宽、缩短相应时间、最大限度地提高性能、加强安全性和支持全球应用；复杂网络的故障排除。通过 HCIE-Enterprise 认证，将证明您对大型网络有全面深入的了解，掌握面向大型企业网络的技术，并具备独立设计各种企业网络以及使用华为路由交换设备实施设计的能力。拥有通过 HCIE-Enterprise 认证的工程师，意味着大中企业有能力独立完成完整的网络搭建，将企业中所需的语音、无线、云、安全和存储全面地集成到网络之中，并且能满足各种应用对网络的使用需求；能够提供完整的故障排除能力；能根据企业和网络技术的发展，规划企业网络的发展，并提供高安全性、可用性和可靠性。

华为认证协助您打开行业之窗，开启改变之门，屹立在ICT世界的潮头浪尖！

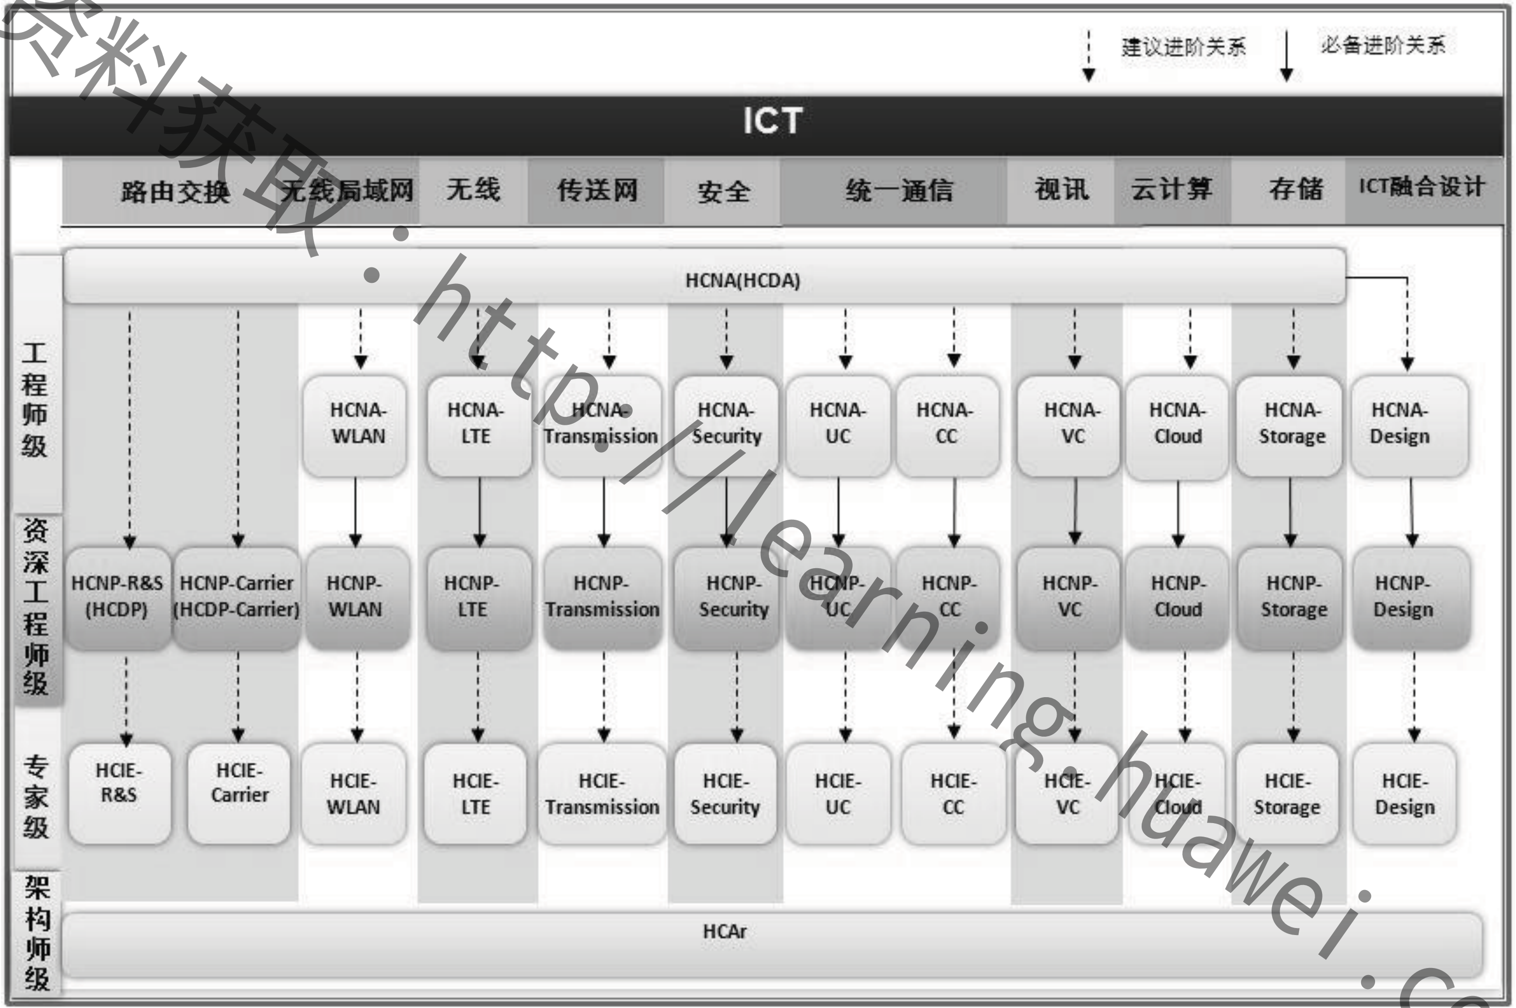
建议进阶关系
必备进阶关系
ICT
路由交换
无线局域网
无线
传送网
安全
统一通信
视讯
云计算
存储
ICT融合设计
工程师级
资深工程师级
专家级
架构师级
HCNA(HCDA)
HCNA-WLAN
HCNA-LTE
HCNA-Transmission
HCNA-Security
HCNA-UC
HCNA-CC
HCNA-VC
HCNA-Cloud
HCNA-Storage
HCNA-Design
HCNP-R&S (HCDP)
HCNP-Carrier (HCDP-Carrier)
HCNP-WLAN
HCNP-LTE
HCNP-Transmission
HCNP-Security
HCNP-UC
HCNP-CC
HCNP-VC
HCNP-Cloud
HCNP-Storage
HCNP-Design
HCIE-R&S
HCIE-Carrier
HCIE-WLAN
HCIE-LTE
HCIE-Transmission
HCIE-Security
HCIE-UC
HCIE-CC
HCIE-VC
HCIE-Cloud
HCIE-Storage
HCIE-Design
HCAr

本书常用图标

路由器

三层交换机

二层交换机

防火墙

网云

以太网线缆

串行线缆

实验环境说明

组网介绍

本实验环境面向准备HCDP-IERN考试的网络工程师，实验设备包括路由器5台，交换机4台，防火墙2台。每套实验环境适用于2名学员同时上机操作。

设备介绍

为了满足HCDP-IERN实验需要，建议每套实验环境采用以下配置：

设备名称、型号与版本的对应关系如下：

设备名称	设备型号	软件版本
R1	AR 2220	Version 5.90 (V200R001C01SPC300)
R2	AR 2220	Version 5.90 (V200R001C01SPC300)
R3	AR 2220	Version 5.90 (V200R001C01SPC300)
R4	AR 1220	Version 5.90 (V200R001C01SPC300)
R5	AR 1220	Version 5.90 (V200R001C01SPC300)
S1	S5700-28C-EI-24S	Version 5.70 (V100R006C00SPC800)
S2	S5700-28C-EI-24S	Version 5.70 (V100R006C00SPC800)
S3	S3700-28TP-EI-AC	Version 5.70 (V100R006C00SPC800)
S4	S3700-28TP-EI-AC	Version 5.70 (V100R006C00SPC800)
FW1	USG2160	Version 5.30 (V300R001C00SPC700)
FW2	USG2160	Version 5.30 (V300R001C00SPC700)

目录

第一章 RIP协议特性与配置

实验 1-1 RIPv1&v2

实验目标

- 掌握RIPv2的配置方式
- 掌握RIPv2的汇总方式
- 掌握RIPv2的认证配置方式
- 掌握RIPv2中被动接口的配置及效果
- 掌握RIPv2与RIPv1的兼容配置

拓扑图

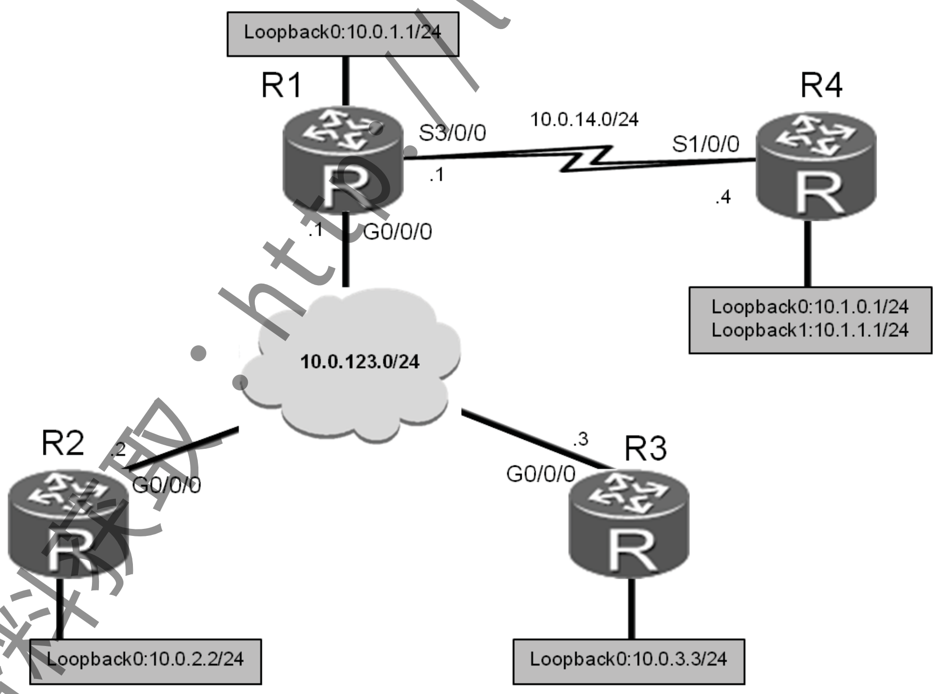

图1-1 RIPv1&v2

场景

R1、R2和R3通过以太交换机相连一个网段10.0.123.0/24，R1和R4之间采用串口线相连，全网使用RIP路由协议。依次完成本实验中的各步骤，以理解RIP协议的配置、路由汇总、认证的配置、RIPv1和RIPv2的兼容配置等知识。

学习任务

步骤一. IP 编址与基本配置

给所有路由器配置IP地址信息。

```
<Huawei>system-view
Enter system view, return user view with Ctrl+Z.
[Huawei]sysname R1
[R1]interface loopback 0
[R1-LoopBack0]ip address 10.0.1.1 24
[R1-LoopBack0]interface Serial 3/0/0
[R1-Serial3/0/0]ip address 10.0.14.1 24
[R1-Serial3/0/0]interface GigabitEthernet 0/0/0
[R1-GigabitEthernet0/0/0]ip address 10.0.123.1 24

<Huawei>system-view
Enter system view, return user view with Ctrl+Z.
[Huawei]sysname R2
[R2]interface GigabitEthernet 0/0/0
[R2-GigabitEthernet0/0/0]ip address 10.0.123.2 24
[R2-GigabitEthernet0/0/0]interface loopback 0
[R2-LoopBack0]ip address 10.0.2.2 24

<Huawei>system-view
Enter system view, return user view with Ctrl+Z.
[Huawei]sysname R3
[R3]interface GigabitEthernet 0/0/0
[R3-GigabitEthernet0/0/0]ip address 10.0.123.3 24
[R3-GigabitEthernet0/0/0]interface loopback 0
[R3-LoopBack0]ip address 10.0.3.3 24

<Huawei>system-view
Enter system view, return user view with Ctrl+Z.
[Huawei]sysname R4
```

```
[R4]interface serial 1/0/0
[R4-Serial1/0/0]ip address 10.0.14.4 24
[R4-Serial1/0/0]interface loopback 0
[R4-LoopBack0]ip address 10.1.0.1 24
[R4-LoopBack0]interface loopback 1
[R4-LoopBack1]ip address 10.1.1.1 24
```

配置完成后，在R1上测试到R2、R3和R4的连通性。

```
[R1]ping 10.0.123.2
  PING 10.0.123.2:  56  data bytes, press CTRL_C to break
    Reply from 10.0.123.2:  bytes=56 Sequence=1 ttl=255 time=3 ms
    Reply from 10.0.123.2:  bytes=56 Sequence=2 ttl=255 time=2 ms
    Reply from 10.0.123.2:  bytes=56 Sequence=3 ttl=255 time=2 ms
    Reply from 10.0.123.2:  bytes=56 Sequence=4 ttl=255 time=2 ms
    Reply from 10.0.123.2:  bytes=56 Sequence=5 ttl=255 time=2 ms

  --- 10.0.123.2 ping statistics ---
    5 packet(s) transmitted
    5 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 2/2/3 ms

[R1]ping 10.0.123.3
  PING 10.0.123.3:  56  data bytes, press CTRL_C to break
    Reply from 10.0.123.3:  bytes=56 Sequence=1 ttl=255 time=3 ms
    Reply from 10.0.123.3:  bytes=56 Sequence=2 ttl=255 time=2 ms
    Reply from 10.0.123.3:  bytes=56 Sequence=3 ttl=255 time=2 ms
    Reply from 10.0.123.3:  bytes=56 Sequence=4 ttl=255 time=2 ms
    Reply from 10.0.123.3:  bytes=56 Sequence=5 ttl=255 time=2 ms

  --- 10.0.123.3 ping statistics ---
    5 packet(s) transmitted
    5 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 2/2/3 ms

[R1]ping 10.0.14.4
  PING 10.0.14.4:  56  data bytes, press CTRL_C to break
    Reply from 10.0.14.4:  bytes=56 Sequence=1 ttl=255 time=42 ms
    Reply from 10.0.14.4:  bytes=56 Sequence=2 ttl=255 time=37 ms
    Reply from 10.0.14.4:  bytes=56 Sequence=3 ttl=255 time=37 ms
    Reply from 10.0.14.4:  bytes=56 Sequence=4 ttl=255 time=37 ms
    Reply from 10.0.14.4:  bytes=56 Sequence=5 ttl=255 time=37 ms
```

```
--- 10.0.14.4 ping statistics ---
  5 packet(s) transmitted
  5 packet(s) received
  0.00% packet loss
  round-trip min/avg/max = 37/38/42 ms
```

步骤二. 配置 RIPv2 协议

所有路由器配置RIP协议。

配置时使用RIP的版本二，关闭自动汇总。

配置RIP时，可以配置进程号码，此号码在本实验无特殊意义，实验中使用进程号码1。

```
[R1]rip 1
[R1-rip-1]version 2
[R1-rip-1]network 10.0.0.0
[R1-rip-1]undo summary

[R2]rip 1
[R2-rip-1]version 2
[R2-rip-1]network 10.0.0.0
[R2-rip-1]undo summary

[R3]rip 1
[R3-rip-1]version 2
[R3-rip-1]network 10.0.0.0
[R3-rip-1]undo summary

[R4]rip 1
[R4-rip-1]version 2
[R4-rip-1]network 10.0.0.0
[R4-rip-1]undo summary
```

配置完成后，观察R2的路由表，并测试R2到R4的连通性。

```
[R2]display ip routing-table protocol rip
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Public routing table : RIP
        Destinations : 5        Routes : 5
```

```
RIP routing table status : <Active>
        Destinations : 5       Routes : 5

Destination/Mask    Proto   Pre  Cost      Flags NextHop        Interface

      10.0.1.0/24  RIP    100  1          D   10.0.123.1   GigabitEthernet0/0/0
      10.0.3.0/24  RIP    100  1          D   10.0.123.3   GigabitEthernet0/0/0
     10.0.14.0/24  RIP    100  1          D   10.0.123.1   GigabitEthernet0/0/0
      10.1.0.0/24  RIP    100  2          D   10.0.123.1   GigabitEthernet0/0/0
      10.1.1.0/24  RIP    100  2          D   10.0.123.1   GigabitEthernet0/0/0

RIP routing table status : <Inactive>
        Destinations : 0       Routes : 0

[R2]ping 10.0.14.4
  PING 10.0.14.4: 56  data bytes, press CTRL_C to break
    Reply from 10.0.14.4: bytes=56 Sequence=1 ttl=255 time=37 ms
    Reply from 10.0.14.4: bytes=56 Sequence=2 ttl=255 time=32 ms
    Reply from 10.0.14.4: bytes=56 Sequence=3 ttl=255 time=32 ms
    Reply from 10.0.14.4: bytes=56 Sequence=4 ttl=255 time=32 ms
    Reply from 10.0.14.4: bytes=56 Sequence=5 ttl=255 time=32 ms

  --- 10.0.14.4 ping statistics ---
    5 packet(s) transmitted
    5 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 32/33/37 ms
```

从以上输出可以看到，R2从G0/0/0接口共学习到5条RIP路由，并且此时R2可以可以与R4通讯。

步骤三. 在 R4 上进行手动汇总

在R4配置对路由10.1.0.0/24和10.1.1.0/24进行手动汇总，汇总为路由10.1.0.0/23。

```
[R4]interface Serial 1/0/0
[R4-Serial1/0/0]rip summary-address 10.1.0.0 255.255.254.0
```

配置完成后，在R2上观察路由表。

```
[R2]display ip routing-table protocol rip
```

```
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Public routing table : RIP
        Destinations : 4       Routes : 4

RIP routing table status : <Active>
        Destinations : 4       Routes : 4

Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface

      10.0.1.0/24  RIP     100  1           D   10.0.123.1   GigabitEthernet0/0/0
      10.0.3.0/24  RIP     100  1           D   10.0.123.3   GigabitEthernet0/0/0
     10.0.14.0/24  RIP     100  1           D   10.0.123.1   GigabitEthernet0/0/0
      10.1.0.0/23  RIP     100  2           D   10.0.123.1   GigabitEthernet0/0/0

RIP routing table status : <Inactive>
        Destinations : 0       Routes : 0
```

我们发现只出现了一条汇总路由10.1.0.0/23，而10.1.0.0/24和10.1.1.0/24两条明细路由消失了。

删除R4的loopback1接口IP地址，即删除了10.1.1.0/24这个网段，观察R2上路由表是否发生变化。

```
[R4]interface loopback 1
[R4-LoopBack1]undo ip address

[R2]display ip routing-table protocol rip
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Public routing table : RIP
        Destinations : 4       Routes : 4

RIP routing table status : <Active>
        Destinations : 4       Routes : 4

Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface

      10.0.1.0/24  RIP     100  1           D   10.0.123.1   GigabitEthernet0/0/0
      10.0.3.0/24  RIP     100  1           D   10.0.123.3   GigabitEthernet0/0/0
     10.0.14.0/24  RIP     100  1           D   10.0.123.1   GigabitEthernet0/0/0
      10.1.0.0/23  RIP     100  2           D   10.0.123.1   GigabitEthernet0/0/0

RIP routing table status : <Inactive>
```

```
        Destinations : 0        Routes : 0
```

我们发现R2上的路由表条目没有变化，依然显示为一条汇总路由，但是如果将R4上的Loopback0接口IP地址也删除，则R2上的汇总路由会消失，可自行验证。

步骤四. 在 10.0.123.0/24 网段启用 RIPv2 的明文认证

在R1、R2和R3的接口G0/0/0上配置RIP认证，认证模式为明文，密码为huawei。

```
[R1]interface GigabitEthernet 0/0/0
[R1-GigabitEthernet0/0/0]rip authentication-mode simple huawei

[R2]interface GigabitEthernet 0/0/0
[R2-GigabitEthernet0/0/0]rip authentication-mode simple huawei

[R2]display ip routing-table protocol rip
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Public routing table : RIP
        Destinations : 3        Routes : 3

RIP routing table status : <Active>
        Destinations : 3        Routes : 3

Destination/Mask    Proto   Pre  Cost      Flags NextHop        Interface

      10.0.1.0/24  RIP     100  1          D   10.0.123.1   GigabitEthernet0/0/0
      10.0.14.0/24  RIP     100  1          D   10.0.123.1   GigabitEthernet0/0/0
      10.1.0.0/23  RIP     100  2          D   10.0.123.1   GigabitEthernet0/0/0

RIP routing table status : <Inactive>
        Destinations : 0        Routes : 0
```

从R2的路由表里可以看出，R2从R1学习到了两条路由路由，而没有从R3学习到路由。

此时R3未配置RIP认证。观查R3路由学习情况。

```
[R3]display ip routing-table
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Routing Tables。 Public
```

```
        Destinations : 10      Routes : 10

Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface

      10.0.3.0/24  Direct  0    0           D   10.0.3.3        LoopBack0
      10.0.3.1/32  Direct  0    0           D   127.0.0.1       InLoopBack0
    10.0.3.255/32  Direct  0    0           D   127.0.0.1       InLoopBack0
    10.0.123.0/24  Direct  0    0           D   10.0.123.3   GigabitEthernet0/0/0
    10.0.123.3/32  Direct  0    0           D   127.0.0.1       InLoopBack0
  10.0.123.255/32  Direct  0    0           D   127.0.0.1       InLoopBack0
     127.0.0.0/8   Direct  0    0           D   127.0.0.1       InLoopBack0
     127.0.0.1/32  Direct  0    0           D   127.0.0.1       InLoopBack0
127.255.255.255/32  Direct  0    0           D   127.0.0.1       InLoopBack0
255.255.255.255/32  Direct  0    0           D   127.0.0.1       InLoopBack0
```

由输出信息，发现R3的路由表里面没有任何从RIP学习到的路由，因为R3此时与R1和R2的RIP认证不匹配，无法学习到路由。

可以使用**debug**命令观察R3对收到的路由信息处理的情况。

```
<R3>terminal debugging
<R3>debugging rip 1 packet
Nov  9 2011 11: 39: 08.180.1+00: 00 R3 RM/6/RMDEBUG:  6:  11698:  RIP 1:  Receive
response from 10.0.123.2 on GigabitEthernet0/0/0
Nov  9 2011 11: 39: 08.180.2+00: 00 R3 RM/6/RMDEBUG:  6:  11709:  Packet:  Version
2, Cmd response, Length 44
Nov  9 2011 11: 39: 08.180.3+00: 00 R3 RM/6/RMDEBUG:  6:  11833:  Authentication-mode
- Simple:  huawei
Nov  9 2011 11: 39: 08.180.4+00: 00 R3 RM/6/RMDEBUG:  6:  11777:  Dest 10.0.2.0/24,
Nexthop 0.0.0.0, Cost 1, Tag 0
Nov  9 2011 11: 39: 08.180.5+00: 00 R3 RM/3/RMDEBUG:  6:  10855:  RIP 1:  Authentication
failure
Nov  9 2011 11: 39: 08.180.6+00: 00 R3 RM/6/RMDEBUG:  6:  1662:  RIP 1:  Process message
failed
Nov  9 2011 11: 39: 14.800.1+00: 00 R3 RM/6/RMDEBUG:  6:  11689:  RIP 1:  Sending
response on interface GigabitEthernet0/0/0 from 10.0.123.3 to 224.0.0.9
<R3>undo debug all
```

我们发现R3从GE0/0/0接口收到了来自其它路由器的RIP路由信息包，但是由于认证信息不匹配导致认证失败。

使用**debug**命令观察出R2对收到R3与R1发过来的路由信息处理的情况。

```
<R2>terminal debugging
Info:  Current terminal debugging is on.
```

```
<R2>debugging rip 1 packet
Nov  9 2011 11: 39: 16.260.1+00: 00 R2 RM/6/RMDEBUG:  6:  11698:  RIP 1:  Receive
response from 10.0.123.1 on GigabitEthernet0/0/0
Nov  9 2011 11: 39: 16.260.2+00: 00 R2 RM/6/RMDEBUG:  6:  11709:  Packet:  Version
2, Cmd response, Length 84
Nov  9 2011 11: 39: 16.260.3+00: 00 R2 RM/6/RMDEBUG:  6:  11833:  Authentication-mode
- Simple:  huawei
Nov  9 2011 11: 39: 16.260.4+00: 00 R2 RM/6/RMDEBUG:  6:  11777:  Dest 10.0.1.0/24,
Nexthop 0.0.0.0, Cost 1, Tag 0
Nov  9 2011 11: 39: 16.260.5+00: 00 R2 RM/6/RMDEBUG:  6:  11777:  Dest 10.0.14.0/24,
Nexthop 0.0.0.0, Cost 1, Tag 0
Nov  9 2011 11: 39: 16.260.6+00: 00 R2 RM/6/RMDEBUG:  6:  11777:  Dest 10.1.0.0/16,
Nexthop 0.0.0.0, Cost 2, Tag 0
Nov  9 2011 11: 39: 23.940.1+00: 00 R2 RM/6/RMDEBUG:  6:  11689:  RIP 1:  Sending
response on interface GigabitEthernet0/0/0 from 10.0.123.2 to 224.0.0.9
Nov  9 2011 11: 39: 23.940.2+00: 00 R2 RM/6/RMDEBUG:  6:  11709:  Packet:  Version
2, Cmd response, Length 44
Nov  9 2011 11: 39: 23.940.3+00: 00 R2 RM/6/RMDEBUG:  6:  11833:  Authentication-mode
- Simple:  huawei
Nov  9 2011 11: 39: 23.940.4+00: 00 R2 RM/6/RMDEBUG:  6:  11777:  Dest 10.0.2.0/24,
Nexthop 0.0.0.0, Cost 1, Tag 0
<R2>undo debug all
Info:  All possible debugging has been turned off
```

R3配置RIP的明文认证。

```
[R3]interface GigabitEthernet 0/0/0
[R3-GigabitEthernet0/0/0]rip authentication-mode simple huawei
```

查看R3的路由表。由于此时R3与R1和R2的认证信息匹配，我们发现R3可以从其它路由器学到RIP路由。

```
[R3]display ip routing-table protocol rip
Route Flags: R - relay, D - download to fib
----------------------------------------------------------------------------
Public routing table : RIP
        Destinations : 4       Routes : 4

RIP routing table status : <Active>
        Destinations : 4       Routes : 4

Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface

      10.0.1.0/24  RIP     100  1           D   10.0.123.1   GigabitEthernet0/0/0
```

```
      10.0.2.0/24  RIP    100  1          D   10.0.123.2   GigabitEthernet0/0/0
     10.0.14.0/24  RIP    100  1          D   10.0.123.1   GigabitEthernet0/0/0
      10.1.0.0/23  RIP    100  2          D   10.0.123.1   GigabitEthernet0/0/0

RIP routing table status : <Inactive>
        Destinations : 0       Routes : 0
```

在R3上测试与其他网段的连通性。

```
[R3]ping 10.1.0.1
  PING 10.1.0.1:  56  data bytes, press CTRL_C to break
    Reply from 10.1.0.1:  bytes=56 Sequence=1 ttl=254 time=42 ms
    Reply from 10.1.0.1:  bytes=56 Sequence=2 ttl=254 time=38 ms
    Reply from 10.1.0.1:  bytes=56 Sequence=3 ttl=254 time=38 ms
    Reply from 10.1.0.1:  bytes=56 Sequence=4 ttl=254 time=38 ms
    Reply from 10.1.0.1:  bytes=56 Sequence=5 ttl=254 time=38 ms

  --- 10.1.0.1 ping statistics ---
    5 packet(s) transmitted
    5 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 38/38/42 ms
```

步骤五. 在 10.0.14.0 网段启用 RIPv2 的 MD5 认证

首先在R1上配置RIPv2的MD5认证。使用密码为huawei。

```
[R1]interface Serial 3/0/0
[R1-Serial3/0/0]rip authentication-mode md5 usual huawei
```

在R4配置MD5认证之前，我们观察R4的路由表。由于认证不匹配，发现没有任何RIP路由。另外我们也可以使用**debug**命令观察RIP信息学习及认证失败的情况。

```
[R4]display ip routing-table
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Routing Tables: Public
        Destinations : 14       Routes : 14

Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface

     10.0.14.0/24  Direct  0    0           D   10.0.14.4       Serial1/0/0
```

```
     10.0.14.1/32  Direct  0    0           D   10.0.14.1       Serial1/0/0
     10.0.14.4/32  Direct  0    0           D   127.0.0.1       InLoopBack0
   10.0.14.255/32  Direct  0    0           D   127.0.0.1       InLoopBack0
      10.1.0.0/24  Direct  0    0           D   10.1.0.1        LoopBack0
      10.1.0.1/32  Direct  0    0           D   127.0.0.1       InLoopBack0
    10.1.0.255/32  Direct  0    0           D   127.0.0.1       InLoopBack0
      10.1.1.0/24  Direct  0    0           D   10.1.1.1        LoopBack1
      10.1.1.1/32  Direct  0    0           D   127.0.0.1       InLoopBack0
    10.1.1.255/32  Direct  0    0           D   127.0.0.1       InLoopBack0
     127.0.0.0/8   Direct  0    0           D   127.0.0.1       InLoopBack0
     127.0.0.1/32  Direct  0    0           D   127.0.0.1       InLoopBack0
127.255.255.255/32  Direct  0    0           D   127.0.0.1       InLoopBack0
255.255.255.255/32  Direct  0    0           D   127.0.0.1       InLoopBack0

<R4>terminal debugging
Info:  Current terminal debugging is on.
<R4>debugging rip 1 packet
Nov  9 2011 14: 51: 02.810.1+00: 00 R4 RM/6/RMDEBUG:  6:  11698:  RIP 1:  Receive
response from 10.0.14.1 on Serial1/0/0
Nov  9 2011 14: 51: 02.810.2+00: 00 R4 RM/6/RMDEBUG:  6:  11709:  Packet:  Version
2, Cmd response, Length 128
Nov  9 2011 14: 51: 02.810.3+00: 00 R4 RM/6/RMDEBUG:  6:  11869:  Authentication-mode
- MD5 Digest:  f8e1dc30.deb83bb7.1b8ce4e5.0e92e4a1
Nov  9 2011 14: 51: 02.810.4+00: 00 R4 RM/6/RMDEBUG:  6:  11901:  Sequence:  00011728
Nov  9 2011 14: 51: 02.810.5+00: 00 R4 RM/6/RMDEBUG:  6:  11777:  Dest 10.0.1.0/24,
Nexthop 0.0.0.0, Cost 1, Tag 0
Nov  9 2011 14: 51: 02.810.6+00: 00 R4 RM/6/RMDEBUG:  6:  11777:  Dest 10.0.2.0/24,
Nexthop 0.0.0.0, Cost 2, Tag 0
Nov  9 2011 14: 51: 02.810.7+00: 00 R4 RM/6/RMDEBUG:  6:  11777:  Dest 10.0.3.0/24,
Nexthop 0.0.0.0, Cost 2, Tag 0
Nov  9 2011 14: 51: 02.810.8+00: 00 R4 RM/6/RMDEBUG:  6:  11777:  Dest 10.0.14.0/24,
Nexthop 0.0.0.0, Cost 1, Tag 0
Nov  9 2011 14: 51: 02.810.9+00: 00 R4 RM/6/RMDEBUG:  6:  11777:  Dest 10.0.123.0/24,
Nexthop 0.0.0.0, Cost 1, Tag 0
Nov  9 2011 14: 51: 02.810.10+00: 00 R4 RM/3/RMDEBUG:  6:  10855:  RIP 1:
Authentication failure
Nov  9 2011 14: 51: 02.810.11+00: 00 R4 RM/6/RMDEBUG:  6:  1662:  RIP 1:  Process
message failed
```

R4配置RIP的MD5认证。

```
[R4]interface Serial 1/0/0
[R4-Serial1/0/0]rip authentication-mode md5 usual huawei
```

配置完成后，在R4上查看路由信息，我们发现R4又学习到了其它路由器发送来的路由条目。

```
[R4]display ip routing-table protocol rip
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Public routing table : RIP
       Destinations : 4       Routes : 4

RIP routing table status : <Active>
       Destinations : 4       Routes : 4

Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface

     10.0.1.0/24  RIP     100  1           D   10.0.14.1       Serial1/0/0
     10.0.2.0/24  RIP     100  2           D   10.0.14.1       Serial1/0/0
     10.0.3.0/24  RIP     100  2           D   10.0.14.1       Serial1/0/0
    10.0.123.0/24  RIP     100  1           D   10.0.14.1       Serial1/0/0

RIP routing table status : <Inactive>
       Destinations : 0       Routes : 0
```

在R4上测试网络的连通性。

```
[R4]ping 10.0.3.3
  PING 10.0.3.3: 56  data bytes, press CTRL_C to break
    Reply from 10.0.3.3: bytes=56 Sequence=1 ttl=254 time=46 ms
    Reply from 10.0.3.3: bytes=56 Sequence=2 ttl=254 time=30 ms
    Reply from 10.0.3.3: bytes=56 Sequence=3 ttl=254 time=30 ms
    Reply from 10.0.3.3: bytes=56 Sequence=4 ttl=254 time=30 ms
    Reply from 10.0.3.3: bytes=56 Sequence=5 ttl=254 time=48 ms

  --- 10.0.3.3 ping statistics ---
    5 packet(s) transmitted
    5 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 30/36/48 ms

[R4]ping 10.0.2.2
  PING 10.0.2.2: 56  data bytes, press CTRL_C to break
    Reply from 10.0.2.2: bytes=56 Sequence=1 ttl=254 time=33 ms
    Reply from 10.0.2.2: bytes=56 Sequence=2 ttl=254 time=30 ms
    Reply from 10.0.2.2: bytes=56 Sequence=3 ttl=254 time=30 ms
    Reply from 10.0.2.2: bytes=56 Sequence=4 ttl=254 time=30 ms
```

```
  Reply from 10.0.2.2:  bytes=56 Sequence=5 ttl=254 time=30 ms

 --- 10.0.2.2 ping statistics ---
   5 packet(s) transmitted
   5 packet(s) received
   0.00% packet loss
   round-trip min/avg/max = 30/30/33 ms
```

步骤六. 配置 RIP 的抑制接口

实际的网络应用中，基于某些安全原因，在RIP协议配置后，我们需要禁止RIP在某些接口上发送RIP更新。将在R4的S1/0/0接口定义为抑制接口。

```
[R4]rip 1
[R4-rip-1]silent-interface Serial 1/0/0
```

配置完成后关闭R4的S1/0/0端口，然后再打开。使路由器重新学习路由信息，然后观察R4路由信息学习情况。

```
[R4]interface Serial 1/0/0
[R4-Serial1/0/0]shutdown
[R4-Serial1/0/0]undo shutdown
```

操作完成后，稍等片刻，待路由信息学习完整后，观察R4的路由表。

```
[R4]display ip routing-table protocol rip
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Public routing table : RIP
       Destinations : 4        Routes : 4

RIP routing table status : <Active>
       Destinations : 4        Routes : 4

Destination/Mask    Proto   Pre  Cost      Flags NextHop        Interface

      10.0.1.0/24  RIP     100  1           D   10.0.14.1      Serial1/0/0
      10.0.2.0/24  RIP     100  2           D   10.0.14.1      Serial1/0/0
      10.0.3.0/24  RIP     100  2           D   10.0.14.1      Serial1/0/0
    10.0.123.0/24  RIP     100  1           D   10.0.14.1      Serial1/0/0

RIP routing table status : <Inactive>
       Destinations : 0        Routes : 0
```

我们发现R4的路由表没有变化，依然能学习到全网的路由条目。

接下来查看R1和R3的路由表，对比之前的路由表，观察发生的变化。

```
[R1]display ip routing-table protocol rip
Route Flags: R - relay, D - download to fib
----------------------------------------------------------------------------
Public routing table : RIP
       Destinations : 2       Routes : 2

RIP routing table status : <Active>
       Destinations : 2       Routes : 2

Destination/Mask    Proto   Pre  Cost      Flags NextHop        Interface

     10.0.2.0/24  RIP    100  1          D   10.0.123.2   GigabitEthernet0/0/0
     10.0.3.0/24  RIP    100  1          D   10.0.123.3   GigabitEthernet0/0/0

RIP routing table status : <Inactive>
       Destinations : 0       Routes : 0

[R3]display ip routing-table protocol rip
Route Flags: R - relay, D - download to fib
----------------------------------------------------------------------------
Public routing table : RIP
       Destinations : 3       Routes : 3

RIP routing table status : <Active>
       Destinations : 3       Routes : 3

Destination/Mask    Proto   Pre  Cost      Flags NextHop        Interface

     10.0.1.0/24  RIP    100  1          D   10.0.123.1   GigabitEthernet0/0/0
     10.0.2.0/24  RIP    100  1          D   10.0.123.2   GigabitEthernet0/0/0
    10.0.14.0/24  RIP    100  1          D   10.0.123.1   GigabitEthernet0/0/0

RIP routing table status : <Inactive>
       Destinations : 0       Routes : 0
```

如输出信息，发现R1和R3的路由表里都不再有10.1.0.0/23这个网段的路由。因为我们将R4的接口S1/0/0设置成为了被动接口，此时R4就只会接受网络中的路由，而不再发布自身的路由信息了，所以R1、R3和R2都不能学习到R4上的路由条目了。

步骤七. 配置 RIPv1 和 RIPv2 的兼容

在R2和R3上将RIP协议的版本改成1。

```
[R2]rip 1
[R2-rip-1]version 1

[R3]rip 1
[R3-rip-1]version 1
```

关闭并重新开启R1的接口G0/0/0以使路由重新计算。

```
[R4]interface GigabitEthernet 0/0/0
[R4-GigabitEthernet0/0/0]shutdown
[R4-GigabitEthernet0/0/0]undo shutdown
```

观察R1计算完成后的各路由器的路由表。

```
[R1]display ip routing-table protocol rip
[R1]
```

从以上R1的路由表里看出，没有任何RIP路由条目，因为一方面R4配置了抑制端口，停止向R1发送路由更新；另一方面由于R2和R3发送过来的路由信息为RIP的版本1，R1运行RIP版本2，而无法接受R2和R3发来的路由信息。可以通过debug命令查看路由信息的收发情况。

```
<R1>terminal debugging
Info: Current terminal debugging is on.
<R1>debugging rip 1 packet
Nov 10 2011 20:47:53.190.1+00:00 R1 RM/6/RMDEBUG: 6: 11689: RIP 1: Sending response
on interface GigabitEthernet0/0/0 from 10.0.123.1 to 224.0.0.9
Nov 10 2011 20:47:53.190.2+00:00 R1 RM/6/RMDEBUG: 6: 11709: Packet: Version 2,
Cmd response, Length 84
Nov 10 2011 20:48:13.540.1+00:00 R1 RM/6/RMDEBUG: 6: 11698: RIP 1: Receive response
from 10.0.123.3 on GigabitEthernet0/0/0
Nov 10 2011 20:48:13.540.2+00:00 R1 RM/6/RMDEBUG: 6: 11709: Packet: Version 1,
Cmd response, Length 24
Nov 10 2011 20:48:13.540.3+00:00 R1 RM/6/RMDEBUG: 6: 11758: Dest 10.0.3.0, Cost
1
Nov 10 2011 20:48:13.540.4+00:00 R1 RM/6/RMDEBUG: 6: 2410: RIP 1: Ignoring packet.
This version is not configured.
Nov 10 2011 20:48:13.980.1+00:00 R1 RM/6/RMDEBUG: 6: 11698: RIP 1: Receive response
from 10.0.123.2 on GigabitEthernet0/0/0
Nov 10 2011 20:48:13.980.2+00:00 R1 RM/6/RMDEBUG: 6: 11709: Packet: Version 1,
```

```
Cmd response, Length 24
Nov 10 2011 20:48:13.980.4+00:00 R1 RM/6/RMDEBUG: 6: 2410: RIP 1: Ignoring packet.
This version is not configured.
<R1>undo debugging all
Info: All possible debugging has been turned off
```

从输出可以看到，R1发送的是版本2的信息，而从R2和R3收到的是版本1的信息。R1因为版本不匹配忽略了这些版本1的信息。

同样R2和R3也因为版本不匹配而学习不到R1上的RIP路由条目，可以使用debug命令查看R2和R3上的路由传播情况。

实际的网络环境中，会存在网络RIPv1和v2共存的情况。RIPv2设计时，考虑到了兼容问题，在恰当的配置下，可以允许RIPv1路由器和RIPv2路由器相互交换路由信息。

在R1的G0/0/0接口下配置在该接口下以版本1的模式运行，这样从该接口下收到版本1的信息时也就能正确识别。

```
[R1]interface GigabitEthernet 0/0/0
[R1-GigabitEthernet0/0/0]rip version 1
[R1-GigabitEthernet0/0/0]display this
[V200R001C00SPC200]
#
interface GigabitEthernet0/0/0
 ip address 10.0.123.1 255.255.255.0
 rip authentication-mode simple huawei
 rip version 1
#
Return
```

然后我们查看R1、R2和R3的路由表，我们发现双方都可以学习到对方的路由条目了。

```
[R1]display ip routing-table protocol rip
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Public routing table : RIP
        Destinations : 2       Routes : 2

RIP routing table status : <Active>
        Destinations : 2       Routes : 2

Destination/Mask    Proto   Pre  Cost      Flags NextHop        Interface

      10.0.2.0/24  RIP    100  1          D   10.0.123.2   GigabitEthernet0/0/0
```

```
      10.0.3.0/24  RIP    100  1          D   10.0.123.3   GigabitEthernet0/0/0

RIP routing table status : <Inactive>
        Destinations : 0      Routes : 0

[R2]display ip routing-table protocol rip
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Public routing table : RIP
        Destinations : 3      Routes : 3

RIP routing table status : <Active>
        Destinations : 3      Routes : 3

Destination/Mask    Proto   Pre  Cost     Flags NextHop        Interface

      10.0.1.0/24  RIP    100  1          D   10.0.123.1   GigabitEthernet0/0/0
      10.0.3.0/24  RIP    100  1          D   10.0.123.3   GigabitEthernet0/0/0
     10.0.14.0/24  RIP    100  1          D   10.0.123.1   GigabitEthernet0/0/0

RIP routing table status : <Inactive>
        Destinations : 0      Routes : 0

[R3]display ip routing-table protocol rip
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Public routing table : RIP
        Destinations : 3      Routes : 3

RIP routing table status : <Active>
        Destinations : 3      Routes : 3

Destination/Mask    Proto   Pre  Cost     Flags NextHop        Interface

      10.0.1.0/24  RIP    100  1          D   10.0.123.1   GigabitEthernet0/0/0
      10.0.2.0/24  RIP    100  1          D   10.0.123.2   GigabitEthernet0/0/0
     10.0.14.0/24  RIP    100  1          D   10.0.123.1   GigabitEthernet0/0/0

RIP routing table status : <Inactive>
        Destinations : 0      Routes : 0
```

配置文件参考

```
[R1]display current-configuration
[V200R001C00SPC200]
#
 sysname R1
#
interface Serial3/0/0
 link-protocol ppp
 ip address 10.0.14.1 255.255.255.0
 rip authentication-mode md5 usual gg^dP=F.[>=H)H2[EInB~.2#
interface GigabitEthernet0/0/0
 ip address 10.0.123.1 255.255.255.0
 rip authentication-mode simple huawei
 rip version 1
#
interface LoopBack0
 ip address 10.0.1.1 255.255.255.0
#
rip 1
 undo summary
 version 2
 network 10.0.0.0
#
Return
```

```
[R2]display current-configuration
[V200R001C00SPC200]
#
 sysname R2
#
interface GigabitEthernet0/0/0
 ip address 10.0.123.2 255.255.255.0
 rip authentication-mode simple huawei
#
interface LoopBack0
 ip address 10.0.2.2 255.255.255.0
#
rip 1
 version 1
 network 10.0.0.0
#
Return
```

```
[R3]display current-configuration
[V200R001C00SPC200]
#
 sysname R3
#
interface GigabitEthernet0/0/0
 ip address 10.0.123.3 255.255.255.0
 rip authentication-mode simple huawei
#
interface LoopBack0
 ip address 10.0.3.3 255.255.255.0
#
rip 1
 version 1
 network 10.0.0.0
#
Return
```

```
[R4]display current-configuration
[V200R001C00SPC200]
#
 sysname R4
#
aaa
 authentication-scheme default
 authorization-scheme default
 accounting-scheme default
 domain default
 domain default_admin
 local-user admin password simple admin
 local-user admin service-type http
#
interface Serial1/0/0
 link-protocol ppp
 ip address 10.0.14.4 255.255.255.0
 rip authentication-mode md5 usual gg^dP=F.[>=H)H2[EInB~.2#
 rip summary-address 10.1.0.0 255.255.254.0
#
interface LoopBack0
 ip address 10.1.0.1 255.255.255.0
#
interface LoopBack1
```

```
 ip address 10.1.1.1 255.255.255.0
#
rip 1
 undo summary
 version 2
 network 10.0.0.0
 silent-interface Serial1/0/0
#
return
```

第二章 OSPF协议特性与配置

实验 2-1 OSPF 单区域

学习目的

- 掌握单区域OSPF的配置方法
- 掌握OSPF区域认证的配置方法
- 了解OSPF在多路访问网络邻居关系建立的过程
- 理解OSPF对Loopback接口所连接网络的掩码发布的形式
- 掌握对OSPF接口代价值进行修改的方法
- 掌握OSPF中Silent-interface的配置方法
- 掌握使用Display查看OSPF各种状态的方法
- 掌握使用Debug命令查看OSPF邻接关系和进行故障排除的方法

拓扑图

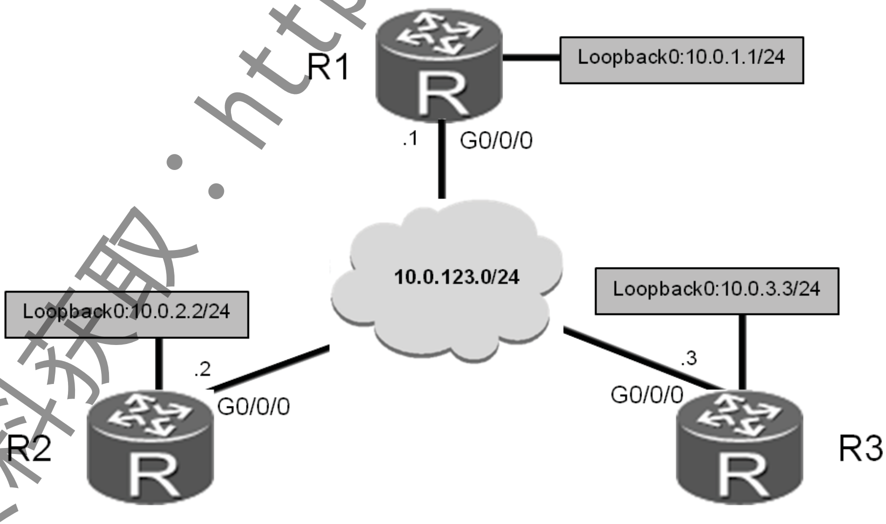

图2-1 OSPF单区域

场景

你是公司的网络管理员。现在公司的网络中有三台ARG3路由器，通过以太网实现相互的连通。在以太网这样的广播式多路访问网络上，可能存在安全隐患，所有你选择采用OSPF区域认证的方法来避免恶意的路由攻击。在部署网络的过程中，出现了网络连通性的问题，你通过使用**display**和**debug**命令进行了故障排除。

学习任务

步骤一. 基础配置与 IP 编址

给R1、R2和R3配置IP地址和掩码。配置时Loopback接口配置掩码为24位，模拟成一个单独的网段。配置完成后，测试直连链路的连通性。

```
<R1>system-view
Enter system view, return user view with Ctrl+Z.
[R1]interface GigabitEthernet 0/0/0
[R1-GigabitEthernet0/0/0]ip address 10.0.123.1 24
[R1-GigabitEthernet0/0/0]quit
[R1]interface LoopBack 0
[R1-LoopBack0]ip address 10.0.1.1 24
[R1-LoopBack0]quit

<R2>system-view
Enter system view, return user view with Ctrl+Z.
[R2]interface GigabitEthernet 0/0/0
[R2-GigabitEthernet0/0/0]ip address 10.0.123.2 24
[R2-GigabitEthernet0/0/0]quit
[R2]interface LoopBack 0
[R2-LoopBack0]ip address 10.0.2.2 24
[R2-LoopBack0]quit

<R3>system-view
Enter system view, return user view with Ctrl+Z.
[R3]interface GigabitEthernet 0/0/0
[R3-GigabitEthernet0/0/0]ip address 10.0.123.3 24
[R3-GigabitEthernet0/0/0]quit
[R3]interface LoopBack 0
[R3-LoopBack0]ip address 10.0.3.3 24
[R3-LoopBack0]quit
```

配置完各接口地址之后验证路由器之间的连通性。

```
[R1]ping -c 1 10.0.123.2
  PING 10.0.123.2: 56  data bytes, press CTRL_C to break
    Reply from 10.0.123.2: bytes=56 Sequence=1 ttl=255 time=2 ms

  --- 10.0.123.2 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 2/2/2 ms

[R1]ping -c 1 10.0.123.3
  PING 10.0.123.3: 56  data bytes, press CTRL_C to break
    Reply from 10.0.123.3: bytes=56 Sequence=1 ttl=255 time=2 ms

  --- 10.0.123.3 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 2/2/2 ms

[R2]ping -c 1 10.0.123.3
  PING 10.0.123.3: 56  data bytes, press CTRL_C to break
    Reply from 10.0.123.3: bytes=56 Sequence=1 ttl=255 time=2 ms

  --- 10.0.123.3 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 2/2/2 ms
```

步骤二.　配置单区域的 OSPF

配置单区域OSPF。所有路由器属于区域0，配置使用OSPF进程1。同时配置区域认证，使用密码“huawei ”。在区域中，华为的设备支持使用明文或MD5值进行认证，在这里，我们仅使用明文进行认证。

注意在使用network命令时，通配符掩码使用0.0.0.0。为了保证路由器的Router ID稳定，我们在启动OSPF进程时使用**router-id**参数静态指定路由器的Router ID。

```
[R1]ospf 1 router-id 10.0.1.1
```

```
[R1-ospf-1]area 0
[R1-ospf-1-area-0.0.0.0]network 10.0.123.1 0.0.0.0
[R1-ospf-1-area-0.0.0.0]network 10.0.1.1 0.0.0.0
[R1-ospf-1-area-0.0.0.0]authentication-mode simple plain huawei

[R2]ospf 1 router-id 10.0.2.2
[R2-ospf-1]area 0
[R2-ospf-1-area-0.0.0.0]network 10.0.123.2 0.0.0.0
[R2-ospf-1-area-0.0.0.0]network 10.0.2.2 0.0.0.0
[R2-ospf-1-area-0.0.0.0]authentication-mode simple plain huawei

[R3]ospf 1 router-id 10.0.3.3
[R3-ospf-1]area 0
[R3-ospf-1-area-0.0.0.0]network 10.0.123.3 0.0.0.0
[R3-ospf-1-area-0.0.0.0]network 10.0.3.3 0.0.0.0
[R3-ospf-1-area-0.0.0.0]authentication-mode simple plain huawei
```

配置完成后，查看设备的路由表，并测试全网的连通性。我们首先在R1上查看路由表。

```
[R1]display ip routing-table
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Routing Tables: Public
        Destinations : 12       Routes : 12

Destination/Mask    Proto  Pre Cost    Flags NextHop        Interface

      10.0.1.0/24   Direct  0   0        D  10.0.1.1        LoopBack0
      10.0.1.1/32   Direct  0   0        D  127.0.0.1       InLoopBack0
    10.0.1.255/32   Direct  0   0        D  127.0.0.1       InLoopBack0
      10.0.2.2/32   OSPF    10  1        D  10.0.123.2      GigabitEthernet0/0/0
      10.0.3.3/32   OSPF    10  1        D  10.0.123.3      GigabitEthernet0/0/0
    10.0.123.0/24   Direct  0   0        D  10.0.123.1      GigabitEthernet0/0/0
    10.0.123.1/32   Direct  0   0        D  127.0.0.1       InLoopBack0
  10.0.123.255/32   Direct  0   0        D  127.0.0.1       InLoopBack0
     127.0.0.0/8    Direct  0   0        D  127.0.0.1       InLoopBack0
     127.0.0.1/32   Direct  0   0        D  127.0.0.1       InLoopBack0
127.255.255.255/32  Direct  0   0        D  127.0.0.1       InLoopBack0
255.255.255.255/32  Direct  0   0        D  127.0.0.1       InLoopBack0
```

从输出中我们可以看到R1从OSPF学习到了2条路由，10.0.2.2/32和10.0.3.3/32，下一跳分别是10.0.123.2和10.0.123.3。然后分别检查从R1到达R2及R3的Loopback地址的连通性。

```
[R1]ping -c 1 10.0.2.2
  PING 10.0.2.2: 56  data bytes, press CTRL_C to break
    Reply from 10.0.2.2: bytes=56 Sequence=1 ttl=255 time=3 ms

  --- 10.0.2.2 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 3/3/3 ms

[R1]ping -c 1 10.0.3.3
  PING 10.0.3.3: 56  data bytes, press CTRL_C to break
    Reply from 10.0.3.3: bytes=56 Sequence=1 ttl=255 time=2 ms

  --- 10.0.3.3 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 2/2/2 ms
```

使用display ospf brief命令查看路由器运行的基本OSPF信息。

```
[R1]display ospf brief

         OSPF Process 1 with Router ID 10.0.1.1
                 OSPF Protocol Information

 RouterID: 10.0.1.1        Border Router:
 Multi-VPN-Instance is not enabled
 Global DS-TE Mode: Non-Standard IETF Mode
 Graceful-restart capability: disabled
 Helper support capability  : not configured
 Spf-schedule-interval: max 10000ms, start 500ms, hold 1000ms
 Default ASE parameters: Metric: 1 Tag: 1 Type: 2
 Route Preference: 10
 ASE Route Preference: 150
 SPF Computation Count: 18
 RFC 1583 Compatible
 Retransmission limitation is disabled
 Area Count: 1   Nssa Area Count: 0
 ExChange/Loading Neighbors: 0

 Area: 0.0.0.0
 Authtype: Simple   Area flag: Normal
```

```
SPF scheduled Count: 18
ExChange/Loading Neighbors: 0

Interface: 10.0.123.1 (GigabitEthernet0/0/0)
Cost: 1       State: DROther    Type: Broadcast    MTU: 1500
Priority: 1
Designated Router: 10.0.123.2
Backup Designated Router: 10.0.123.3
Timers: Hello 10 , Dead 40 , Poll  120 , Retransmit 5 , Transmit Delay 1

Interface: 10.0.1.1 (LoopBack0)
Cost: 0       State: P-2-P     Type: P2P      MTU: 1500
Timers: Hello 10 , Dead 40 , Poll  120 , Retransmit 5 , Transmit Delay 1
```

从上面的输出中我们可以看到区域0开启了明文认证(Authtype: Simple)，共有两个接口参加了OSPF的运行 :GigabitEthernet0/0/0和LoopBack0。其中，GigabitEthernet0/0/0为广播型网络（Broadcast），开销（Cost）为1，优先级（Priority）为1，R1自己的角色为DROther，后面列出了该网络上的DR（10.0.123.2）和BDR（10.0.123.3）。另外一个运行OSPF的接口LoopBack0的网络类型为P2P。

使用display ospf peer brief命令查看路由器的OSPF邻居关系建立情况。

```
[R1]display ospf peer brief

       OSPF Process 1 with Router ID 10.0.1.1
              Peer Statistic Information
 ----------------------------------------------------------------------------
 Area Id          Interface                     Neighbor id      State
 0.0.0.0          GigabitEthernet0/0/0            10.0.2.2         Full
 0.0.0.0          GigabitEthernet0/0/0            10.0.3.3         Full
 ----------------------------------------------------------------------------
```

从上面的输出中我们可以看到在区域0.0.0.0中，R1有两个邻居，邻居的Router ID分别为10.0.2.2和10.0.3.3，他们之间的状态为Full。

使用display ospf lsdb命令查看路由器的OSPF数据库信息。

```
[R1]dis ospf lsdb

       OSPF Process 1 with Router ID 10.0.1.1
             Link State Database

                   Area: 0.0.0.0
 Type      LinkState ID      AdvRouter          Age  Len   Sequence   Metric
```

```
Router    10.0.3.3       10.0.3.3        98  48    80000011     1
Router    10.0.2.2       10.0.2.2        98  48    80000016     1
Router    10.0.1.1       10.0.1.1        96  48    80000013     1
Network   10.0.123.2     10.0.2.2        99  36    8000000F     0
```

在这里我们一共可以看到4条LSA，前3条为第一类LSA，分别由R1、R2和R3产生，我们可以通过AdvRouter判断该LSA是由哪台路由器生成的。第四条为第二类LSA，是由一个网段的DR产生的。在这里，R2是10.0.123.0/24这个网段的DR，所以我们可以看到这条LSA的AdvRouter为10.0.2.2。

```
[R1]display ospf lsdb router self-originate

       OSPF Process 1 with Router ID 10.0.1.1
                    Area: 0.0.0.0
              Link State Database


 Type      : Router
 Ls id     : 10.0.1.1
 Adv rtr   : 10.0.1.1
 Ls age    : 680
 Len       : 48
 Options   :  E
 seq#      : 80000013
 chksum    : 0x7787
 Link count: 2
  * Link ID: 10.0.123.2
    Data   : 10.0.123.1
    Link Type: TransNet
    Metric : 1
  * Link ID: 10.0.1.1
    Data   : 255.255.255.255
    Link Type: StubNet
    Metric : 0
    Priority : Medium
```

上面的输出是R1产生的Router LSA的详细信息 ,我们可以看到这条LSA一共描述了2个网络，第一个网络为三台路由器的互联网段，链路类型为TransNet，可以看到Link ID为DR的端口地址10.0.123.2，Data为该网段上本地端口的IP地址10.0.123.1；另一个网络为Loopback接口所在网段，链路类型为StubNet，Link ID和Data分别是该Stub网段的IP地址和掩码。

```
[R2]dis ospf lsdb network self-originate

       OSPF Process 1 with Router ID 10.0.2.2
```

```
                    Area: 0.0.0.0
              Link State Database

  Type      : Network
  Ls id     : 10.0.123.2
  Adv rtr    : 10.0.2.2
  Ls age     : 1369
  Len       : 36
  Options    :  E
  seq#      : 8000000f
  chksum     : 0xa7e
  Net mask  : 255.255.255.0
  Priority  : Low
    Attached Router    10.0.2.2
    Attached Router    10.0.1.1
    Attached Router    10.0.3.3
```

上面的输出是R2产生的Network LSA的详细信息，我们可以看到第二类LSA描述了DR所在网段的邻居信息。

步骤三. 观察路由器在以太网上邻接关系的建立过程

首先查看在10.0.123.0/24网段，OSPF邻居关系中DR和BDR选举的情况，并分析为什么会这样？以及是否所有人在做这个实验时，结果都是一样的。

我们首先查看在10.0.123.0/24网段，OSPF邻居关系中DR和BDR选举的情况。从下面的输出中，我们可以得知现在该网段的DR的接口IP为10.0.123.2，BDR的接口IP为10.0.123.3。

```
[R1]display ospf peer

       OSPF Process 1 with Router ID 10.0.1.1
             Neighbors

 Area 0.0.0.0 interface 10.0.123.1(GigabitEthernet0/0/0)'s neighbors
 Router ID: 10.0.2.2        Address: 10.0.123.2
   State: Full  Mode:Nbr is  Master  Priority: 1
   DR: 10.0.123.2  BDR: 10.0.123.3  MTU: 0
   Dead timer due in 40  sec
   Retrans timer interval: 5
   Neighbor is up for 00:32:27
   Authentication Sequence: [ 0 ]
```

```
 Router ID: 10.0.3.3        Address: 10.0.123.3
   State: Full  Mode:Nbr is  Master  Priority: 1
   DR: 10.0.123.2  BDR: 10.0.123.3  MTU: 0
   Dead timer due in 33  sec
   Retrans timer interval: 3
   Neighbor is up for 00:32:28
   Authentication Sequence: [ 0 ]
```

有可能每个人得实验结果输出不一样。因为在OSPF中，DR的选举不是抢占的，即网络中存在DR或BDR时，新进入网络的路由器不能抢占DR或BDR的角色。在这个网络中，先启动OSPF进程或先接入该网络的路由器成为了该网段上的DR，其他路由器成为的BDR或DROther。

当DR发生故障后，BDR就会接替DR的位置，我们在实验中可以通过重置OSPF进程的方法来观察DR角色的改变，在这里，我们重置R2的OSPF进程。

```
<R2>reset ospf process
Warning: The OSPF process will be reset. Continue? [Y/N]:y

[R2]display ospf peer

        OSPF Process 1 with Router ID 10.0.2.2
                Neighbors

 Area 0.0.0.0 interface 10.0.123.2(GigabitEthernet0/0/0)'s neighbors
 Router ID: 10.0.1.1        Address: 10.0.123.1
   State: Full  Mode:Nbr is  Slave  Priority: 1
   DR: 10.0.123.3  BDR: 10.0.123.1  MTU: 0
   Dead timer due in 33  sec
   Retrans timer interval: 4
   Neighbor is up for 00:00:35
   Authentication Sequence: [ 0 ]

 Router ID: 10.0.3.3        Address: 10.0.123.3
   State: Full  Mode:Nbr is  Master  Priority: 1
   DR: 10.0.123.3  BDR: 10.0.123.1  MTU: 0
   Dead timer due in 33  sec
   Retrans timer interval: 5
   Neighbor is up for 00:00:35
   Authentication Sequence: [ 0 ]
```

当重置R2的OSPF进程以后，原来该网络上的BDR 10.0.123.3成为了新的DR，原来的DROther 10.0.123.1成为了新的BDR。

下面我们关闭R1、R2与R3的G0/0/0接口，使用命令**debugging ospf 1 event**准备查看OSPF邻接关系建立的具体过程。然后尽量同时打开R1、R2与R3的G0/0/0接口。观察在广播式多路访问网络上邻居状态的变化过程和DR和BDR选举的过程。

```
<R1>debugging ospf 1 event
<R1>terminal debugging
[R1]int GigabitEthernet 0/0/0
[R1-GigabitEthernet0/0/0]shut
[R1-GigabitEthernet0/0/0]undo shut
```

在R2和R3上进行相同的操作，查看R3的debug信息。由于所有路由器默认的接口优先级都是1，所以在DR选举的时候会参考路由器的Router ID，在这三台路由器中，R3的Router ID是最大的，所以R3成为了该网段上的DR。

```
[R3-GigabitEthernet0/0/0]undo shutdown
Nov 22 2011 18:41:50.990.3+00:00 R3 RM/6/RMDEBUG:
 FileID: 0xd017802c Line: 1268 Level: 0x20
  OSPF 1: Intf 10.0.123.3 Rcv InterfaceUp State Down -> Waiting.
Nov 22 2011 18:41:50.990.4+00:00 R3 RM/6/RMDEBUG:
 FileID: 0xd017802c Line: 1382 Level: 0x20
  OSPF 1 Send Hello Interface Up on 10.0.123.3
Nov 22 2011 18:41:57.470.1+00:00 R3 RM/6/RMDEBUG:
 FileID: 0xd017802d Line: 1132 Level: 0x20
  OSPF 1: Nbr 10.0.123.1 Rcv HelloReceived State Down -> Init.
Nov 22 2011 18:41:57.480.1+00:00 R3 RM/6/RMDEBUG:
 FileID: 0xd017802d Line: 1728 Level: 0x20
  OSPF 1: Nbr 10.0.123.1 Rcv 2WayReceived State Init -> 2Way.
Nov 22 2011 18:41:59.510.3+00:00 R3 RM/6/RMDEBUG:
 FileID: 0xd017802d Line: 1132 Level: 0x20
  OSPF 1: Nbr 10.0.123.2 Rcv HelloReceived State Down -> Init.
Nov 22 2011 18:41:59.510.4+00:00 R3 RM/6/RMDEBUG:
 FileID: 0xd017802d Line: 1728 Level: 0x20
  OSPF 1: Nbr 10.0.123.2 Rcv 2WayReceived State Init -> 2Way.
Nov 22 2011 18:42:28.350.4+00:00 R3 RM/6/RMDEBUG:
 FileID: 0xd017802d Line: 1728 Level: 0x20
  OSPF 1: Nbr 10.0.123.1 Rcv AdjOk? State 2Way -> ExStart.
Nov 22 2011 18:42:28.350.5+00:00 R3 RM/6/RMDEBUG:
 FileID: 0xd017802d Line: 1728 Level: 0x20
  OSPF 1: Nbr 10.0.123.2 Rcv AdjOk? State 2Way -> ExStart.
Nov 22 2011 18:42:28.350.6+00:00 R3 RM/6/RMDEBUG:
 FileID: 0xd017802c Line: 2045 Level: 0x20
  OSPF 1 Send Hello Interface State Changed on 10.0.123.3
Nov 22 2011 18:42:28.350.7+00:00 R3 RM/6/RMDEBUG:
```

```
FileID: 0xd017802c Line: 2056 Level: 0x20
 OSPF 1: Intf 10.0.123.3 Rcv WaitTimer State Waiting -> DR.
```

当刚打开接口时，接口状态由Down变为Waiting，此时路由器开始交互Hello数据包，等待约40秒以后，R3的接口由Waiting变为DR。

步骤四. 配置 OSPF 中 Loopback 接口的网络类型

观察R1的路由表，关注这两条路由：10.0.2.2/32和10.0.3.3/32。

```
[R1]display ip routing-table
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Routing Tables: Public
        Destinations : 12       Routes : 12

Destination/Mask    Proto  Pre Cost    Flags NextHop        Interface

      10.0.1.0/24   Direct  0    0       D  10.0.1.1       LoopBack0
      10.0.1.1/32   Direct  0    0       D  127.0.0.1      InLoopBack0
    10.0.1.255/32   Direct  0    0       D  127.0.0.1      InLoopBack0
      10.0.2.2/32   OSPF   10    1       D  10.0.123.2     GigabitEthernet0/0/0
      10.0.3.3/32   OSPF   10    1       D  10.0.123.3     GigabitEthernet0/0/0
    10.0.123.0/24   Direct  0    0       D  10.0.123.1     GigabitEthernet0/0/0
    10.0.123.1/32   Direct  0    0       D  127.0.0.1      InLoopBack0
  10.0.123.255/32   Direct  0    0       D  127.0.0.1      InLoopBack0
     127.0.0.0/8    Direct  0    0       D  127.0.0.1      InLoopBack0
     127.0.0.1/32   Direct  0    0       D  127.0.0.1      InLoopBack0
127.255.255.255/32  Direct  0    0       D  127.0.0.1      InLoopBack0
255.255.255.255/32  Direct  0    0       D  127.0.0.1      InLoopBack0
```

在配置R2和R3的Loopback接口地址时，使用的掩码是24位，分析为什么这里为什么路由表中显示的是32位掩码的路由？

使用命令**display ospf interface LoopBack 0 verbose**查看OSPF在Loopback 0接口运行的状态信息。

```
[R1]display ospf interface LoopBack 0 verbose

        OSPF Process 1 with Router ID 10.0.1.1
                Interfaces


 Interface: 10.0.1.1 (LoopBack0)
```

```
 Cost: 0       State: P-2-P      Type: P2P      MTU: 1500
 Timers: Hello 10 , Dead 40 , Poll  120 , Retransmit 5 , Transmit Delay 1
  IO Statistics
             Type        Input     Output
            Hello           0          0
   DB Description           0          0
   Link-State Req           0          0
Link-State Update           0          0
   Link-State Ack           0          0
 ALLSPF GROUP
 OpaqueId: 0   PrevState: Down
```

可以看到对于Loopback接口，OSPF知道该网段只可能有一个IP地址，所以发布的路由的子网掩码是32位的。

修改R2的Loopback0接口的网络类型为Broadcast，OSPF在发布这个接口的网络信息时，就会使用24位掩码进行发布。

```
[R2]interface LoopBack 0
[R2-LoopBack0]ospf network-type broadcast
```

这时我们看到R2发布的Loopback 0地址的路由子网掩码为24位。

```
[R1]display ip routing-table
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Routing Tables: Public
        Destinations : 12       Routes : 12

Destination/Mask    Proto  Pre Cost    Flags NextHop        Interface

      10.0.1.0/24   Direct  0   0        D   10.0.1.1       LoopBack0
      10.0.1.1/32   Direct  0   0        D   127.0.0.1       InLoopBack0
    10.0.1.255/32   Direct  0   0        D   127.0.0.1       InLoopBack0
      10.0.2.0/24   OSPF    10  1        D   10.0.123.2     GigabitEthernet0/0/0
      10.0.3.3/32   OSPF    10  1        D   10.0.123.3     GigabitEthernet0/0/0
    10.0.123.0/24   Direct  0   0        D   10.0.123.1     GigabitEthernet0/0/0
    10.0.123.1/32   Direct  0   0        D   127.0.0.1       InLoopBack0
  10.0.123.255/32   Direct  0   0        D   127.0.0.1       InLoopBack0
     127.0.0.0/8    Direct  0   0        D   127.0.0.1       InLoopBack0
     127.0.0.1/32   Direct  0   0        D   127.0.0.1       InLoopBack0
127.255.255.255/32  Direct  0   0        D   127.0.0.1       InLoopBack0
255.255.255.255/32  Direct  0   0        D   127.0.0.1       InLoopBack0
```

使用命令**display ospf interface LoopBack 0 verbose**查看Loopback接口的运行状态可以看到，该接口网络类型为Broadcast。

```
[R2]display ospf interface LoopBack 0 verbose

        OSPF Process 1 with Router ID 10.0.2.2
               Interfaces



 Interface: 10.0.2.2 (LoopBack0)
 Cost: 0      State: DR       Type: Broadcast    MTU: 1500
 Priority: 1
 Designated Router: 10.0.2.2
 Backup Designated Router: 0.0.0.0
 Timers: Hello 10 , Dead 40 , Poll  120 , Retransmit 5 , Transmit Delay 1
  IO Statistics
           Type       Input     Output
          Hello         0         0
   DB Description       0         0
   Link-State Req       0         0
Link-State Update       0         0
   Link-State Ack       0         0
 ALLSPF GROUP
 ALLDR GROUP
 OpaqueId: 0   PrevState: Waiting
```

步骤五. 修改接口的 OSPF 代价值

首先在R1上查看R1到达R3的Loopback0接口路由的代价值，我们可以看到到达10.0.3.3/32的代价值为1。

```
[R1]display ip routing-table
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Routing Tables: Public
        Destinations : 12       Routes : 12

Destination/Mask    Proto  Pre Cost    Flags NextHop        Interface

      10.0.1.0/24   Direct  0    0        D   10.0.1.1        LoopBack0
      10.0.1.1/32   Direct  0    0        D   127.0.0.1       InLoopBack0
    10.0.1.255/32   Direct  0    0        D   127.0.0.1       InLoopBack0
```

```
     10.0.2.0/24   OSPF    10   1        D   10.0.123.2     GigabitEthernet0/0/0
     10.0.3.3/32   OSPF    10   1        D   10.0.123.3     GigabitEthernet0/0/0
   10.0.123.0/24   Direct  0    0        D   10.0.123.1     GigabitEthernet0/0/0
   10.0.123.1/32   Direct  0    0        D   127.0.0.1       InLoopBack0
 10.0.123.255/32   Direct  0    0        D   127.0.0.1       InLoopBack0
    127.0.0.0/8    Direct  0    0        D   127.0.0.1       InLoopBack0
    127.0.0.1/32   Direct  0    0        D   127.0.0.1       InLoopBack0
127.255.255.255/32 Direct  0    0        D   127.0.0.1       InLoopBack0
255.255.255.255/32 Direct  0    0        D   127.0.0.1       InLoopBack0
```

修改R1的G0/0/0接口代价值为20，修改R3的G0/0/0接口代价值为10。

```
[R1]interface GigabitEthernet 0/0/0
[R1-GigabitEthernet0/0/0]ospf cost 20

[R3]interface GigabitEthernet 0/0/0
[R3-GigabitEthernet0/0/0]ospf cost 10
```

重新查看R1到达R3的Loopback0接口路由的代价值，可以看到，到达10.0.3.3/32的代价值为20。

```
[R1]display ip routing-table
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Routing Tables: Public
        Destinations : 12       Routes : 12

Destination/Mask    Proto  Pre  Cost     Flags NextHop         Interface

     10.0.1.0/24   Direct  0    0        D   10.0.1.1        LoopBack0
     10.0.1.1/32   Direct  0    0        D   127.0.0.1       InLoopBack0
   10.0.1.255/32   Direct  0    0        D   127.0.0.1       InLoopBack0
     10.0.2.0/24   OSPF    10   20       D   10.0.123.2     GigabitEthernet0/0/0
     10.0.3.3/32   OSPF    10   20       D   10.0.123.3     GigabitEthernet0/0/0
   10.0.123.0/24   Direct  0    0        D   10.0.123.1     GigabitEthernet0/0/0
   10.0.123.1/32   Direct  0    0        D   127.0.0.1       InLoopBack0
 10.0.123.255/32   Direct  0    0        D   127.0.0.1       InLoopBack0
    127.0.0.0/8    Direct  0    0        D   127.0.0.1       InLoopBack0
    127.0.0.1/32   Direct  0    0        D   127.0.0.1       InLoopBack0
127.255.255.255/32 Direct  0    0        D   127.0.0.1       InLoopBack0
255.255.255.255/32 Direct  0    0        D   127.0.0.1       InLoopBack0
```

在R3上查看10.0.1.1/32的代价值，可以看到值为10。

```
[R3]display ip routing-table
```

```
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Routing Tables: Public
        Destinations : 12       Routes : 12

Destination/Mask    Proto  Pre  Cost   Flags NextHop         Interface

      10.0.1.1/32   OSPF   10   10       D   10.0.123.1      GigabitEthernet0/0/0
      10.0.2.0/24   OSPF   10   10       D   10.0.123.2      GigabitEthernet0/0/0
      10.0.3.0/24   Direct  0   0        D   10.0.3.3        LoopBack0
      10.0.3.3/32   Direct  0   0        D   127.0.0.1       InLoopBack0
    10.0.3.255/32   Direct  0   0        D   127.0.0.1       InLoopBack0
    10.0.123.0/24   Direct  0   0        D   10.0.123.3      GigabitEthernet0/0/0
    10.0.123.3/32   Direct  0   0        D   127.0.0.1       InLoopBack0
  10.0.123.255/32   Direct  0   0        D   127.0.0.1       InLoopBack0
     127.0.0.0/8    Direct  0   0        D   127.0.0.1       InLoopBack0
     127.0.0.1/32   Direct  0   0        D   127.0.0.1       InLoopBack0
127.255.255.255/32  Direct  0   0        D   127.0.0.1       InLoopBack0
255.255.255.255/32  Direct  0   0        D   127.0.0.1       InLoopBack0
```

步骤六.　配置 OSPF 的 Silent-interface

配置R1的G0/0/0接口为Silent-interface。

```
[R1]ospf 1
[R1-ospf-1]silent-interface GigabitEthernet 0/0/0
```

查看R1的邻居关系建立和路由表学习情况可发现，路由表中从OSPF学习到的路由条目消失了。

```
[R1]display ip routing-table
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Routing Tables: Public
        Destinations : 12       Routes : 12

Destination/Mask    Proto  Pre  Cost   Flags NextHop         Interface

      10.0.1.0/24   Direct  0   0        D   10.0.1.1        LoopBack0
      10.0.1.1/32   Direct  0   0        D   127.0.0.1       InLoopBack0
    10.0.1.255/32   Direct  0   0        D   127.0.0.1       InLoopBack0
    10.0.123.0/24   Direct  0   0        D   10.0.123.1      GigabitEthernet0/0/0
```

```
    10.0.123.1/32  Direct  0    0        D   127.0.0.1       InLoopBack0
  10.0.123.255/32  Direct  0    0        D   127.0.0.1       InLoopBack0
     127.0.0.0/8   Direct  0    0        D   127.0.0.1       InLoopBack0
     127.0.0.1/32  Direct  0    0        D   127.0.0.1       InLoopBack0
127.255.255.255/32 Direct  0    0        D   127.0.0.1       InLoopBack0
255.255.255.255/32 Direct  0    0        D   127.0.0.1       InLoopBack0
```

查看R1的邻居列表可以看到R1和R2、R3之间的邻居关系也消失了。在RIP中将一个接口置为Silent-interface以后，该接口不再发送RIP更新；但在OSPF中，路由器之间需要建立邻居关系之后才会交互路由信息，当一个接口被设置为Silent-interface以后，该接口不再接收或发送Hello包，造成该接口不能和其他路由器形成邻居关系。

```
[R1]display ospf interface GigabitEthernet 0/0/0

        OSPF Process 1 with Router ID 10.0.1.1
                Interfaces


 Interface: 10.0.123.1 (GigabitEthernet0/0/0)
 Cost: 20      State: DR       Type: Broadcast    MTU: 1500
 Priority: 10
 Designated Router: 10.0.123.1
 Backup Designated Router: 0.0.0.0
 Timers: Hello 10 , Dead 40 , Poll  120 , Retransmit 5 , Transmit Delay 1
 Silent interface, No hellos
```

使用**debugging ip packet**调试路由器接收到的数据包观察OSPF报文的接收情况，可以看到路由器接收到一个OSPF报文（protocol = 89）以后将该报文丢弃了。

```
<R1>debugging ip packet
Nov 23 2011 09:51:53.500.1+00:00 R1 IP/7/debug_case:
Receiving, interface = GE0/0/0, version = 4, headlen = 20, tos = 192,
pktlen = 68, pktid = 7272, offset = 0, ttl = 1, protocol = 89,
checksum = 14129, s = 10.0.123.3, d = 224.0.0.5
prompt: Receiving IP packet from GE0/0/0


Nov 23 2011 09:51:53.500.2+00:00 R1 IP/7/debug_case:
Receiving, interface = GigabitEthernet0/0/0, version = 4, headlen = 20, tos =
192,
pktlen = 68, pktid = 7272, offset = 0, ttl = 1, protocol = 89,
checksum = 14129, s = 10.0.123.3, d = 224.0.0.5
```

```
prompt: IP_ProcessByBoard Begin!


Nov 23 2011 09:51:53.500.3+00:00 R1 IP/7/debug_case:
Discarding, interface = GE0/0/0, version = 4, headlen = 20, tos = 192,
pktlen = 68, pktid = 7272, offset = 0, ttl = 1, protocol = 89,
checksum = 14129, s = 10.0.123.3, d = 224.0.0.5
prompt: IP_Distribute: The packet was dropped by security application.
```

恢复R1的G0/0/0接口为默认状态，将三个路由器的Loopback0接口配置为Silent-interface。

```
[R1]ospf 1
[R1-ospf-1]undo silent-interface GigabitEthernet0/0/0
[R1-ospf-1]silent-interface LoopBack 0

[R2]ospf 1
[R2-ospf-1]silent-interface LoopBack 0

[R3]ospf 1
[R3-ospf-1]silent-interface LoopBack 0
```

检查R1的路由表可见，将Loopback设为Silent-interface以后不影响该接口路由的发布。

```
[R1]display ip routing-table
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Routing Tables: Public
        Destinations : 12       Routes : 12

Destination/Mask    Proto  Pre  Cost    Flags NextHop        Interface

      10.0.1.0/24   Direct  0    0        D   10.0.1.1        LoopBack0
      10.0.1.1/32   Direct  0    0        D   127.0.0.1       InLoopBack0
    10.0.1.255/32   Direct  0    0        D   127.0.0.1       InLoopBack0
      10.0.2.2/24   OSPF    10   20       D   10.0.123.2      GigabitEthernet0/0/0
      10.0.3.3/32   OSPF    10   20       D   10.0.123.3      GigabitEthernet0/0/0
    10.0.123.0/24   Direct  0    0        D   10.0.123.1      GigabitEthernet0/0/0
    10.0.123.1/32   Direct  0    0        D   127.0.0.1       InLoopBack0
  10.0.123.255/32   Direct  0    0        D   127.0.0.1       InLoopBack0
     127.0.0.0/8    Direct  0    0        D   127.0.0.1       InLoopBack0
     127.0.0.1/32   Direct  0    0        D   127.0.0.1       InLoopBack0
127.255.255.255/32  Direct  0    0        D   127.0.0.1       InLoopBack0
```

附加实验：思考并验证

为什么在配置OSPF时，使用的通配符掩码是0.0.0.0，实际的配置中，也可以使用通配符掩码0.0.0.255，思考一下，这两种表达形式有什么差异？

分析在实际的网络中，哪些类型的接口应该配置为Silent-interface接口？

最终设备配置

```
<R1>display current-configuration
[V200R001C00SPC200]
#
 sysname R1
#
interface GigabitEthernet0/0/0
 ip address 10.0.123.1 255.255.255.0
 ospf cost 20
#
interface LoopBack0
 ip address 10.0.1.1 255.255.255.0
#
ospf 1 router-id 10.0.1.1
 silent-interface LoopBack0
 area 0.0.0.0
  authentication-mode simple plain huawei
  network 10.0.123.1 0.0.0.0
  network 10.0.1.1 0.0.0.0
#
return
```

```
<R2>display current-configuration
[V200R001C00SPC200]
#
 sysname R2
#
interface GigabitEthernet0/0/0
 ip address 10.0.123.2 255.255.255.0
#
interface LoopBack0
```

```
 ip address 10.0.2.2 255.255.255.0
 ospf network-type broadcast
#
ospf 1 router-id 10.0.2.2
 silent-interface LoopBack0
 area 0.0.0.0
  authentication-mode simple plain huawei
  network 10.0.123.2 0.0.0.0
  network 10.0.2.2 0.0.0.0
#
return

<R3>display current-configuration
[V200R001C00SPC200]
#
 sysname R3
#
interface GigabitEthernet0/0/0
 ip address 10.0.123.3 255.255.255.0
 ospf cost 10
#
interface LoopBack0
 ip address 10.0.3.3 255.255.255.0
#
ospf 1 router-id 10.0.3.3
 silent-interface LoopBack0
 area 0.0.0.0
  authentication-mode simple plain huawei
  network 10.0.123.3 0.0.0.0
  network 10.0.3.3 0.0.0.0
#
return
```

实验 2-2 OSPF 多区域

学习目的

- 掌握OSPF配置指定Router ID的方法
- 掌握多区域OSPF的配置方法
- 掌握OSPF区域之间路由汇总的配置方法
- 掌握OSPF参考带宽的配置方法
- 掌握OSPF引入外部路由的配置方法
- 掌握OSPF引入的外部路由时进行路由汇总的方法
- 掌握向OSPF导入缺省路由的方法
- 掌握对OSPF中各类路由的管理距离的修改方法

拓扑图

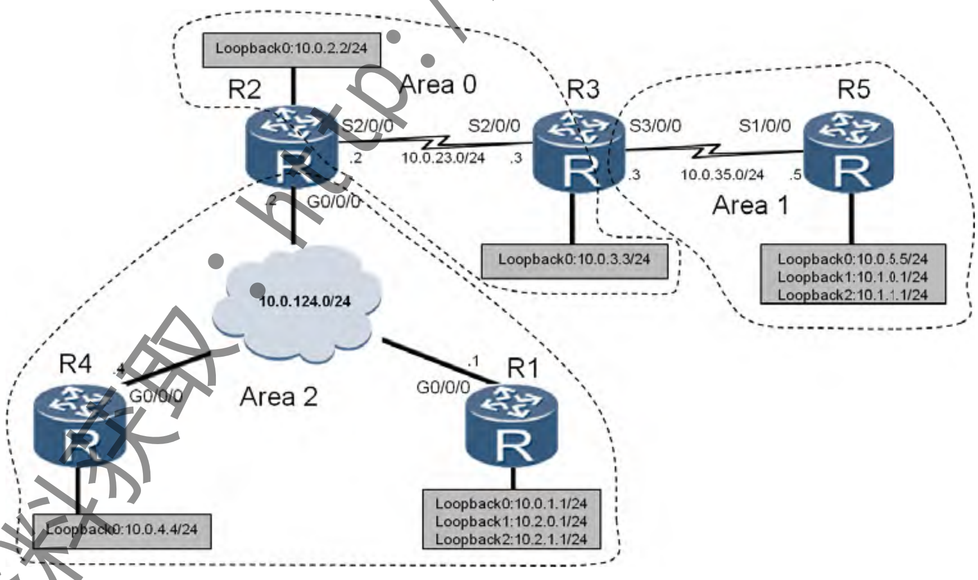

图2-2 OSPF多区域

场景

你是公司的网络管理员。现在公司的网络中有五台ARG3路由器，其中R1、R2和R4在公司总部，通过以太网互联。R3、与R5在公司分部，R3通过专线与公司总部的R2相连，R5与R3之间也通过专项相连。由于网络规模较大，为了控制LSA的洪泛，你设计了多区域的OSPF互联方式。

其中R2与R3的Loopback0接口、互联接口属于区域0 ;R3与R5互联的网段、R5的Loopback0/1/2接口属于区域1；R1、R2与R4互联的网段以及R1、R4的Loopback0接口属于区域2。

同时为了明确设备的Router-ID ,你配置设备使用固定的地址作为Router ID。

为了使路由器进行路由转发时效率更高，在区域的边界你配置了自动汇总。

R1路由器连接到公司以外的网络，你配置将这些OSPF区域之外的路由信息引入到OSPF区域。

R4路由器连接到Internet，你需要配置一条缺省路由，引入到OSPF区域，以便于OSPF区域的所有路由器都知道如何访问Internet。

同时OSPF路由信息中区分了内部路由和外部路由，你修改了OSPF路由信息的优先级信息，以避免潜在的风险。

OSPF中特定路由信息的度量值是将到达目的网络所经过的所有链路的代价值进行累加得到的。而链路的代价值是路由器将接口带宽与参考带宽进行对比得到。参考带宽值为100Mbps，实际接口带宽可能为1000Mbps，而度量值都是整数，所以快速以太网接口和千兆以太网接口的OSPF代价值均为1。为了能够相互区分这些链路，你定义参考带宽值为10Gbps。

在配置设备的同时，出现了一些网络故障，你通过使用display和debug命令进行了故障排除。

学习任务

步骤一.　基础配置与 IP 编址

给所有路由器配置IP地址和掩码。配置时注意所有的Loopback接口配置掩码均为24位，模拟成一个单独的网段。

```
<R1>system-view
Enter system view, return user view with Ctrl+Z.
[R1]interface GigabitEthernet 0/0/0
[R1-GigabitEthernet0/0/0]ip address 10.0.124.1 24
[R1-GigabitEthernet0/0/0]interface LoopBack 0
```

```
[R1-LoopBack0]ip address 10.0.1.1 24
[R1-LoopBack0]interface LoopBack 1
[R1-LoopBack1]ip address 10.2.0.1 24
[R1-LoopBack1]interface LoopBack 2
[R1-LoopBack2]ip address 10.2.1.1 24

<R2>system-view
Enter system view, return user view with Ctrl+Z.
[R2]interface GigabitEthernet 0/0/0
[R2-GigabitEthernet0/0/0]ip address 10.0.124.2 24
[R2-GigabitEthernet0/0/0]interface Serial 2/0/0
[R2-Serial2/0/0]ip address 10.0.23.2 24
[R2-Serial2/0/0]interface LoopBack 0
[R2-LoopBack0]ip address 10.0.2.2 24

<R3>system-view
Enter system view, return user view with Ctrl+Z.
[R3]interface Serial 2/0/0
[R3-Serial2/0/0]ip address 10.0.23.3 24
[R3-Serial2/0/0]interface Serial 3/0/0
[R3-Serial3/0/0]ip address 10.0.35.3 24
[R3-Serial3/0/0]interface LoopBack 0
[R3-LoopBack0]ip address 10.0.3.3 24

<R4>system-view
Enter system view, return user view with Ctrl+Z.
[R4]interface GigabitEthernet 0/0/0
[R4-GigabitEthernet0/0/0]ip address 10.0.124.4 24
[R4-GigabitEthernet0/0/0]interface LoopBack 0
[R4-LoopBack0]ip address 10.0.4.4 24

<R5>system-view
Enter system view, return user view with Ctrl+Z.
[R5]interface Serial 1/0/0
[R5-Serial1/0/0]ip address 10.0.35.5 24
[R5-Serial1/0/0]interface LoopBack 0
[R5-LoopBack0]ip address 10.0.5.5 24
[R5-LoopBack0]interface LoopBack 1
[R5-LoopBack1]ip address 10.1.0.1 24
[R5-LoopBack1]interface LoopBack 2
[R5-LoopBack2]ip address 10.1.1.1 24
```

配置完成后，测试直连链路的连通性。

```
[R2]ping -c 1 10.0.124.1
  PING 10.0.124.1: 56  data bytes, press CTRL_C to break
    Reply from 10.0.124.1: bytes=56 Sequence=1 ttl=255 time=5 ms

  --- 10.0.124.1 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 5/5/5 ms

[R2]ping -c 1 10.0.124.4
  PING 10.0.124.4: 56  data bytes, press CTRL_C to break
    Reply from 10.0.124.4: bytes=56 Sequence=1 ttl=255 time=14 ms

  --- 10.0.124.4 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 14/14/14 ms

[R2]ping -c 1 10.0.23.3
  PING 10.0.23.3: 56  data bytes, press CTRL_C to break
    Reply from 10.0.23.3: bytes=56 Sequence=1 ttl=255 time=41 ms

  --- 10.0.23.3 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 41/41/41 ms

[R3]ping -c 1 10.0.35.5
  PING 10.0.35.5: 56  data bytes, press CTRL_C to break
    Reply from 10.0.35.5: bytes=56 Sequence=1 ttl=255 time=38 ms

  --- 10.0.35.5 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 38/38/38 ms
```

步骤二. 配置多区域 OSPF

为保证OSPF的Router ID稳定，我们通常手工指定路由器的Router ID。有2种方法可以手工指定一台路由器运行OSPF的Router ID，第一种方式是在系统视图下使用**router id**的命令。

```
[R1]router id 10.0.1.1
```

第二种方式是在启动OSPF进程时加上参数**router-id**。

```
[R1]ospf 1 router-id 10.0.1.1
```

当路由器上同时配置了这两条命令以后，路由器最终会选取第二种方式配置的值作为Router ID。如果在一台路由器上需要起多个OSPF进程，且每个OSPF进程的Router ID需要不一样时，我们只能使用第二种方式来指定Router ID。

在R1上配置Loopback 0接口及GigabitEthernet 0/0/0属于区域2。这里我们将所有OSPF区域的Loopback接口，修改其OSPF网络类型为Broadcast类型，以便于OSPF发布Loopback口的真实掩码信息。

```
[R1]ospf 1 router-id 10.0.1.1
[R1-ospf-1]area 2
[R1-ospf-1-area-0.0.0.2]network 10.0.124.1 0.0.0.0
[R1-ospf-1-area-0.0.0.2]quit
[R1-ospf-1]quit
[R1]interface LoopBack 0
[R1-LoopBack0]ospf network-type broadcast
[R1-LoopBack0]quit
```

在R2上配置Loopback 0和Serial 2/0/0接口属于区域0，GigabitEthernet 0/0/0属于区域2。

```
[R2]ospf 1 router-id 10.0.2.2
[R2-ospf-1]area 0
[R2-ospf-1-area-0.0.0.0]network 10.0.23.2 0.0.0.0
[R2-ospf-1-area-0.0.0.0]network 10.0.2.2 0.0.0.0
[R2-ospf-1-area-0.0.0.0]quit
[R2-ospf-1]area 2
[R2-ospf-1-area-0.0.0.2]network 10.0.124.2 0.0.0.0
[R2-ospf-1-area-0.0.0.2]quit
[R2-ospf-1]quit
[R2]interface LoopBack 0
[R2-LoopBack0]ospf network-type broadcast
[R2-LoopBack0]quit
```

在R3上配置Loopback 0和Serial 2/0/0接口属于区域0，Serial 3/0/0属于区域1。

```
[R3]ospf 1 router-id 10.0.3.3
[R3-ospf-1]area 0
[R3-ospf-1-area-0.0.0.0]network 10.0.3.3 0.0.0.0
[R3-ospf-1-area-0.0.0.0]network 10.0.23.3 0.0.0.0
[R3-ospf-1-area-0.0.0.0]quit
[R3-ospf-1]area 1
[R3-ospf-1-area-0.0.0.1]network 10.0.35.3 0.0.0.0
[R3-ospf-1-area-0.0.0.1]quit
[R3-ospf-1]quit
[R3]interface LoopBack 0
[R3-LoopBack0]ospf network-type broadcast
[R3-LoopBack0]quit
```

在R4上配置Loopback 0及GigabitEthernet 0/0/0属于区域2。

```
[R4]ospf 1 router-id 10.0.4.4
[R4-ospf-1]area 2
[R4-ospf-1-area-0.0.0.2]network 10.0.4.4 0.0.0.0
[R4-ospf-1-area-0.0.0.2]network 10.0.124.4 0.0.0.0
[R4-ospf-1-area-0.0.0.2]quit
[R4-ospf-1]quit
[R4]interface LoopBack 0
[R4-LoopBack0]ospf network-type broadcast
[R4-LoopBack0]quit
```

在R5上配置所有的Loopback接口及Serial 1/0/0属于区域1。

```
[R5]ospf 1 router-id 10.0.5.5
[R5-ospf-1]area 1
[R5-ospf-1-area-0.0.0.1]network 10.0.5.5 0.0.0.0
[R5-ospf-1-area-0.0.0.1]network 10.1.0.1 0.0.0.0
[R5-ospf-1-area-0.0.0.1]network 10.1.1.1 0.0.0.0
[R5-ospf-1-area-0.0.0.1]network 10.0.35.5 0.0.0.0
[R5-ospf-1-area-0.0.0.1]quit
[R5-ospf-1]quit
[R5]interface LoopBack 0
[R5-LoopBack0]ospf network-type broadcast
[R5-LoopBack0]quit
[R5]interface LoopBack 1
[R5-LoopBack1]ospf network-type broadcast
[R5-LoopBack1]quit
[R5]interface LoopBack 2
```

```
[R5-LoopBack2]ospf network-type broadcast
[R5-LoopBack2]quit
```

配置完成后，在R1上查看路由表。

```
[R1]display ip routing-table
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Routing Tables: Public
        Destinations : 24      Routes : 24

Destination/Mask    Proto  Pre  Cost    Flags NextHop         Interface

      10.0.1.0/24   Direct  0    0        D   10.0.1.1        LoopBack0
      10.0.1.1/32   Direct  0    0        D   127.0.0.1        InLoopBack0
    10.0.1.255/32   Direct  0    0        D   127.0.0.1        InLoopBack0
      10.0.2.0/24   OSPF    10   1        D   10.0.124.2      GigabitEthernet0/0/0
      10.0.3.0/24   OSPF    10   1563     D   10.0.124.2      GigabitEthernet0/0/0
      10.0.4.0/24   OSPF    10   1        D   10.0.124.4      GigabitEthernet0/0/0
      10.0.5.0/24   OSPF    10   3125     D   10.0.124.2      GigabitEthernet0/0/0
     10.0.23.0/24   OSPF    10   1563     D   10.0.124.2      GigabitEthernet0/0/0
     10.0.35.0/24   OSPF    10   3125     D   10.0.124.2      GigabitEthernet0/0/0
    10.0.124.0/24   Direct  0    0        D   10.0.124.1      GigabitEthernet0/0/0
    10.0.124.1/32   Direct  0    0        D   127.0.0.1        InLoopBack0
  10.0.124.255/32   Direct  0    0        D   127.0.0.1        InLoopBack0
      10.1.0.0/24   OSPF    10   3125     D   10.0.124.2      GigabitEthernet0/0/0
      10.1.1.0/24   OSPF    10   3125     D   10.0.124.2      GigabitEthernet0/0/0
    10.2.0.255/32   Direct  0    0        D   127.0.0.1        InLoopBack0
      10.2.1.0/24   Direct  0    0        D   10.2.1.1        LoopBack2
      10.2.1.1/32   Direct  0    0        D   127.0.0.1        InLoopBack0
    10.2.1.255/32   Direct  0    0        D   127.0.0.1        InLoopBack0
     127.0.0.0/8    Direct  0    0        D   127.0.0.1        InLoopBack0
     127.0.0.1/32   Direct  0    0        D   127.0.0.1        InLoopBack0
127.255.255.255/32  Direct  0    0        D   127.0.0.1        InLoopBack0
255.255.255.255/32  Direct  0    0        D   127.0.0.1        InLoopBack0
```

该路由器上已拥有全网所有的路由条目。

在R1上测试到其他路由器Loopback接口的连通性。

```
[R1]ping -c 1 10.0.2.2
  PING 10.0.2.2: 56  data bytes, press CTRL_C to break
    Reply from 10.0.2.2: bytes=56 Sequence=1 ttl=255 time=3 ms

  --- 10.0.2.2 ping statistics ---
```

```
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 3/3/3 ms

[R1]ping -c 1 10.0.5.5
  PING 10.0.5.5: 56  data bytes, press CTRL_C to break
    Reply from 10.0.5.5: bytes=56 Sequence=1 ttl=253 time=88 ms

  --- 10.0.5.5 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 88/88/88 ms

[R1]ping -c 1 10.0.4.4
  PING 10.0.4.4: 56  data bytes, press CTRL_C to break
    Reply from 10.0.4.4: bytes=56 Sequence=1 ttl=255 time=3 ms

  --- 10.0.4.4 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 3/3/3 ms
```

我们使用**display ospf brief**命令在R2上查看路由器运行的基本OSPF信息。

```
[R2]display ospf brief

         OSPF Process 1 with Router ID 10.0.2.2
                 OSPF Protocol Information

 RouterID: 10.0.2.2          Border Router:  AREA
 Multi-VPN-Instance is not enabled
 Global DS-TE Mode: Non-Standard IETF Mode
 Graceful-restart capability: disabled
 Helper support capability  : not configured
 Spf-schedule-interval: max 10000ms, start 500ms, hold 1000ms
 Default ASE parameters: Metric: 1 Tag: 1 Type: 2
 Route Preference: 10
 ASE Route Preference: 150
 SPF Computation Count: 19
 RFC 1583 Compatible
 Retransmission limitation is disabled
```

```
Area Count: 2   Nssa Area Count: 0
ExChange/Loading Neighbors: 0

Area: 0.0.0.0
Authtype: None   Area flag: Normal
SPF scheduled Count: 19
ExChange/Loading Neighbors: 0

Interface: 10.0.2.2 (LoopBack0)
Cost: 0      State: DR       Type: Broadcast    MTU: 1500
Priority: 1
Designated Router: 10.0.2.2
Backup Designated Router: 0.0.0.0
Timers: Hello 10 , Dead 40 , Poll  120 , Retransmit 5 , Transmit Delay 1

Interface: 10.0.23.2 (Serial2/0/0) --> 10.0.23.3
Cost: 1562    State: P-2-P     Type: P2P      MTU: 1500
Timers: Hello 10 , Dead 40 , Poll  120 , Retransmit 5 , Transmit Delay 1

Area: 0.0.0.2
Authtype: None   Area flag: Normal
SPF scheduled Count: 17
ExChange/Loading Neighbors: 0

Interface: 10.0.124.2 (GigabitEthernet0/0/0)
Cost: 1      State: BDR      Type: Broadcast    MTU: 1500
Priority: 1
Designated Router: 10.0.124.1
Backup Designated Router: 10.0.124.2
Timers: Hello 10 , Dead 40 , Poll  120 , Retransmit 5 , Transmit Delay 1
```

第一行Border Router: AREA 表示该路由器是一台ABR；如果路由器是一台区域内路由器，该值为空；如果路由器是一台ASBR，该值为 AS。

该路由器共有三个接口参加OSPF运算，我们已手工将Loopback 0接口的网络类型修改为Broadcast。Serial2/0/0的封装类型为PPP，所以默认的网络类型为点对点。另外GigabitEthernet 0/0/0连接到区域2，是广播型网络。

我们在R2上使用**display ospf peer brief**命令查看路由器的OSPF邻居关系建立情况。可以看到，在区域0，R2有一个邻居10.0.3.3，在区域2，R2有2个邻居：10.0.1.1和10.0.4.4，R2与他们都形成了邻接关系（Full）。

```
[R2]display ospf peer brief

        OSPF Process 1 with Router ID 10.0.2.2
```

```
              Peer Statistic Information
----------------------------------------------------------------------------
Area Id          Interface                       Neighbor id      State
0.0.0.0          Serial2/0/0                      10.0.3.3         Full
0.0.0.2          GigabitEthernet0/0/0             10.0.1.1         Full
0.0.0.2          GigabitEthernet0/0/0             10.0.4.4         Full
----------------------------------------------------------------------------
```

我们在R2上使用**display ospf lsdb**命令查看路由器的OSPF数据库信息。我们可以发现由于R2是一台ABR，所以在该路由器上维护了2个LSDB，分别用来描述区域0和区域2的路由。

```
[R2]display ospf lsdb

       OSPF Process 1 with Router ID 10.0.2.2
              Link State Database

                       Area: 0.0.0.0
Type      LinkState ID    AdvRouter         Age  Len   Sequence   Metric
Router      10.0.3.3        10.0.3.3          788  60    80000008       0
Router      10.0.2.2        10.0.2.2          869  60    80000008       0
Sum-Net     10.0.35.0       10.0.3.3          846  28    80000002    1562
Sum-Net     10.0.124.0      10.0.2.2         1259  28    80000002       1
Sum-Net     10.1.1.0        10.0.3.3         1565  28    80000001    1562
Sum-Net     10.0.5.0        10.0.3.3         1594  28    80000001    1562
Sum-Net     10.1.0.0        10.0.3.3         1584  28    80000001    1562
Sum-Net     10.0.4.0        10.0.2.2          538  28    80000002       1

                       Area: 0.0.0.2
Type      LinkState ID    AdvRouter         Age  Len   Sequence   Metric
Router      10.0.4.4        10.0.4.4          504  48    80000008       1
Router      10.0.2.2        10.0.2.2          558  36    80000006       1
Router      10.0.1.1        10.0.1.1          568  60    80000011       1
Network     10.0.124.1      10.0.1.1          559  36    80000005       0
Sum-Net     10.0.35.0       10.0.2.2          846  28    80000002    3124
Sum-Net     10.0.3.0        10.0.2.2          830  28    80000002    1562
Sum-Net     10.0.2.0        10.0.2.2         1249  28    80000002       0
Sum-Net     10.1.1.0        10.0.2.2         1565  28    80000001    3124
Sum-Net     10.0.5.0        10.0.2.2         1595  28    80000001    3124
Sum-Net     10.1.0.0        10.0.2.2         1584  28    80000001    3124
Sum-Net     10.0.23.0       10.0.2.2         1261  28    80000002    1562
```

步骤三. 配置 OSPF 区域之间的路由汇总

首先查看R2和R3的OSPF路由表。

```
[R2]display ip routing-table protocol ospf
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Public routing table : OSPF
        Destinations : 8       Routes : 8

OSPF routing table status : <Active>
        Destinations : 8       Routes : 8

Destination/Mask    Proto  Pre  Cost   Flags NextHop         Interface


      10.0.3.0/24  OSPF    10   1562     D   10.0.23.3       Serial2/0/0
      10.0.4.0/24  OSPF    10   1        D   10.0.124.4    GigabitEthernet0/0/0
      10.0.5.0/24  OSPF    10   3124     D   10.0.23.3       Serial2/0/0
     10.0.35.0/24  OSPF    10   3124     D   10.0.23.3       Serial2/0/0
      10.1.0.0/24  OSPF    10   3124     D   10.0.23.3       Serial2/0/0
      10.1.1.0/24  OSPF    10   3124     D   10.0.23.3       Serial2/0/0

OSPF routing table status : <Inactive>
        Destinations : 0       Routes : 0

[R3]display ip routing-table protocol ospf
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Public routing table : OSPF
        Destinations : 8       Routes : 8

OSPF routing table status : <Active>
        Destinations : 8       Routes : 8

Destination/Mask    Proto  Pre  Cost      Flags NextHop         Interface


      10.0.2.0/24   OSPF    10   1562       D   10.0.23.2       Serial2/0/0
      10.0.4.0/24   OSPF    10   1563       D   10.0.23.2       Serial2/0/0
      10.0.5.0/24   OSPF    10   1562       D   10.0.35.5       Serial3/0/0
    10.0.124.0/24   OSPF    10   1563       D   10.0.23.2       Serial2/0/0
      10.1.0.0/24   OSPF    10   1562       D   10.0.35.5       Serial3/0/0
      10.1.1.0/24   OSPF    10   1562       D   10.0.35.5       Serial3/0/0
```

```
OSPF routing table status : <Inactive>
        Destinations : 0       Routes : 0
```

10.1.0.0/24和10.1.1.0/24两条路由信息均以详细条目出现。

对于这样的路由信息，可以进行汇总，再向其他区域发送。一方面减少其他区域的路由表条目，另外一方面还可以减少路由振荡情况的发生。我们可在R3上使用**abr-summary**的命令将R5的Loopback1和Loopback2接口的网段进行汇总发送。

```
[R3]ospf 1
[R3-ospf-1]area 1
[R3-ospf-1-area-0.0.0.1]abr-summary 10.1.0.0 255.255.254.0
```

配置完成后在R3和R2上分别查看汇总路由信息。

```
[R3]display ip routing-table protocol ospf
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Public routing table : OSPF
        Destinations : 8       Routes : 8

OSPF routing table status : <Active>
        Destinations : 8       Routes : 8

Destination/Mask    Proto   Pre  Cost      Flags NextHop        Interface

      10.0.2.0/24   OSPF    10   1562        D   10.0.23.2      Serial2/0/0
      10.0.4.0/24   OSPF    10   1563        D   10.0.23.2      Serial2/0/0
      10.0.5.0/24   OSPF    10   1562        D   10.0.35.5      Serial3/0/0
    10.0.124.0/24   OSPF    10   1563        D   10.0.23.2      Serial2/0/0
      10.1.0.0/24   OSPF    10   1562        D   10.0.35.5      Serial3/0/0
      10.1.1.0/24   OSPF    10   1562        D   10.0.35.5      Serial3/0/0

OSPF routing table status : <Inactive>
        Destinations : 0       Routes : 0

[R2]display ip routing-table protocol ospf
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Public routing table : OSPF
        Destinations : 7       Routes : 7

OSPF routing table status : <Active>
        Destinations : 7       Routes : 7
```

```
Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface

      10.0.3.0/24  OSPF    10   1562      D   10.0.23.3       Serial2/0/0
      10.0.4.0/24  OSPF    10   1         D   10.0.124.4    GigabitEthernet0/0/0
      10.0.5.0/24  OSPF    10   3124      D   10.0.23.3       Serial2/0/0
     10.0.35.0/24  OSPF    10   3124      D   10.0.23.3       Serial2/0/0
      10.1.0.0/23  OSPF    10   3124      D   10.0.23.3       Serial2/0/0

OSPF routing table status : <Inactive>
        Destinations : 0       Routes : 0
```

通过输出可以看到,在R3的路由表中,这2条路由仍以明细路由的形式出现,在R2上,仅存在汇总路由10.1.0.0/23。

配置完成后,测试其他路由器与网络10.1.0.0/24与10.1.1.0/24的连通性。

```
[R1]ping -c 1 10.1.0.1
  PING 10.1.0.1: 56  data bytes, press CTRL_C to break
    Reply from 10.1.0.1: bytes=56 Sequence=1 ttl=253 time=66 ms

  --- 10.1.0.1 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 66/66/66 ms

[R1]ping -c 1 10.1.1.1
  PING 10.1.1.1: 56  data bytes, press CTRL_C to break
    Reply from 10.1.1.1: bytes=56 Sequence=1 ttl=253 time=66 ms

  --- 10.1.1.1 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 66/66/66 ms

[R2]ping -c 1 10.1.0.1
  PING 10.1.0.1: 56  data bytes, press CTRL_C to break
    Reply from 10.1.0.1: bytes=56 Sequence=1 ttl=254 time=69 ms

  --- 10.1.0.1 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
```

```
    0.00% packet loss
    round-trip min/avg/max = 69/69/69 ms

[R3]ping -c 1 10.1.0.1
  PING 10.1.0.1: 56  data bytes, press CTRL_C to break
    Reply from 10.1.0.1: bytes=56 Sequence=1 ttl=255 time=29 ms

  --- 10.1.0.1 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 29/29/29 ms
```

步骤四. 修改 OSPF 的参考带宽值

在实际网络我们可能使用了千兆甚至万兆以太网。但是由于OSPF的默认参考带宽值为100Mbps，并且接口代价值仅为整数，所以OSPF无法在带宽上区分百兆以太网和千兆以太网。

在R2上修改OSPF的参考带宽值为10Gbps。这里，使用命令**bandwidth-reference**进行修改，相应带宽参数值的单位为Mbps。

```
[R2-ospf-1]bandwidth-reference 10000
```

在R2上查看OSPF邻居关系，以及路由信息学习情况，我们可以看到，在路由表中，Cost值已经发生了变化。

```
[R2]display ip routing-table protocol ospf
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Public routing table : OSPF
         Destinations : 7        Routes : 7

OSPF routing table status : <Active>
         Destinations : 7        Routes : 7

Destination/Mask    Proto   Pre  Cost     Flags NextHop         Interface

      10.0.3.0/24   OSPF    10   65535      D   10.0.23.3       Serial2/0/0
      10.0.4.0/24   OSPF    10   10         D   10.0.124.4    GigabitEthernet0/0/0
      10.0.5.0/24   OSPF    10   67097      D   10.0.23.3       Serial2/0/0
     10.0.35.0/24   OSPF    10   67097      D   10.0.23.3       Serial2/0/0
      10.1.0.0/23   OSPF    10   67097      D   10.0.23.3       Serial2/0/0
```

```
OSPF routing table status : <Inactive>
         Destinations : 0        Routes : 0
```

在运行OSPF的多个区域内，OSPF的参考带宽值必须一致，否则OSPF无法正常工作。修改所有路由器的OSPF参考带宽值为10Gbps。

```
[R1-ospf-1]bandwidth-reference 10000

[R3-ospf-1]bandwidth-reference 10000

[R4-ospf-1]bandwidth-reference 10000

[R5-ospf-1]bandwidth-reference 10000
```

在R2上查看邻居列表、路由表，观察OSPF邻居关系以及路由信息是否正常。

```
[R2]display ospf peer brief

         OSPF Process 1 with Router ID 10.0.2.2
                  Peer Statistic Information
 ----------------------------------------------------------------------------
 Area Id          Interface                     Neighbor id      State
 0.0.0.0          Serial2/0/0                     10.0.3.3         Full
 0.0.0.2          GigabitEthernet0/0/0            10.0.1.1         Full
 0.0.0.2          GigabitEthernet0/0/0            10.0.4.4         Full
 ----------------------------------------------------------------------------

[R2]display ip routing-table protocol ospf
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Public routing table : OSPF
         Destinations : 7        Routes : 7

OSPF routing table status : <Active>
         Destinations : 7        Routes : 7

Destination/Mask    Proto   Pre  Cost    Flags NextHop         Interface

      10.0.3.0/24  OSPF    10   65535     D   10.0.23.3       Serial2/0/0
      10.0.4.0/24  OSPF    10   10        D   10.0.124.4    GigabitEthernet0/0/0
      10.0.5.0/24  OSPF    10   131070    D   10.0.23.3       Serial2/0/0
     10.0.35.0/24  OSPF    10   131070    D   10.0.23.3       Serial2/0/0
      10.1.0.0/23  OSPF    10   131070    D   10.0.23.3       Serial2/0/0
```

```
OSPF routing table status : <Inactive>
        Destinations : 0        Routes : 0
```

如上所示，路由信息正常。另外可测试网络的连通性。

步骤五. 配置将直连路由汇总并引入到 OSPF 区域

R1的Loopback1和Loopback2接口不属于OSPF区域。将这两条直连路由引入到OSPF区域，并在R1上执行路由汇总。

```
[R1]ospf
[R1-ospf-1]import-route direct
[R1-ospf-1]asbr-summary 10.2.0.0 255.255.254.0
```

在R1上查看外部路由信息。

```
[R1]display ospf lsdb ase 10.2.0.0

        OSPF Process 1 with Router ID 10.0.1.1
               Link State Database


  Type      : External
  Ls id     : 10.2.0.0
  Adv rtr    : 10.0.1.1
  Ls age    : 293
  Len       : 36
  Options    :  E
  seq#      : 80000001
  chksum    : 0x2b6
  Net mask  : 255.255.254.0
  TOS 0  Metric: 2
  E type    : 2
  Forwarding Address : 0.0.0.0
  Tag       : 1
  Priority  : Low
```

R1通过一条第五类LSA向其他路由器通告了网段10.2.0.0，子网掩码是255.255.254.0。

在其他路由器上查看汇总路由，并测试网络连通性。

```
[R2]display ip routing-table protocol ospf
```

```
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Public routing table : OSPF
        Destinations : 7       Routes : 7

OSPF routing table status : <Active>
        Destinations : 7       Routes : 7

Destination/Mask    Proto   Pre  Cost      Flags NextHop        Interface

      10.0.1.0/24  O_ASE   150  1         D   10.0.124.1    GigabitEthernet0/0/0
      10.0.3.0/24  OSPF    10   65535     D   10.0.23.3      Serial2/0/0
      10.0.4.0/24  OSPF    10   10        D   10.0.124.4    GigabitEthernet0/0/0
      10.0.5.0/24  OSPF    10   131070    D   10.0.23.3      Serial2/0/0
     10.0.35.0/24  OSPF    10   131070    D   10.0.23.3      Serial2/0/0
      10.1.0.0/23  OSPF    10   131070    D   10.0.23.3      Serial2/0/0
      10.2.0.0/23  O_ASE   150  2         D   10.0.124.1    GigabitEthernet0/0/0

OSPF routing table status : <Inactive>
        Destinations : 0       Routes : 0

[R2]ping -c 1 10.2.0.1
  PING 10.2.0.1: 56  data bytes, press CTRL_C to break
    Reply from 10.2.0.1: bytes=56 Sequence=1 ttl=255 time=2 ms

  --- 10.2.0.1 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 2/2/2 ms

[R2]ping -c 1 10.2.1.1
  PING 10.2.1.1: 56  data bytes, press CTRL_C to break
    Reply from 10.2.1.1: bytes=56 Sequence=1 ttl=255 time=2 ms

  --- 10.2.1.1 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 2/2/2 ms
```

在R2上可以看到一条掩码为23位的汇总路由。

将R1的Loopback 2接口删除，查看R2上路由条目变化情况。我们可以看到，

当Loopback 2接口不存在了，汇总路由仍然存在。

```
[R1]undo interface LoopBack 2

[R2]display ip routing-table protocol ospf
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Public routing table : OSPF
        Destinations : 7       Routes : 7

OSPF routing table status : <Active>
        Destinations : 7       Routes : 7

Destination/Mask   Proto   Pre  Cost    Flags NextHop        Interface

      10.0.1.0/24  O_ASE   150  1         D   10.0.124.1   GigabitEthernet0/0/0
      10.0.3.0/24  OSPF    10   65535     D   10.0.23.3      Serial2/0/0
      10.0.4.0/24  OSPF    10   10        D   10.0.124.4   GigabitEthernet0/0/0
      10.0.5.0/24  OSPF    10   131070    D   10.0.23.3      Serial2/0/0
     10.0.35.0/24  OSPF    10   131070    D   10.0.23.3      Serial2/0/0
      10.1.0.0/23  OSPF    10   131070    D   10.0.23.3      Serial2/0/0
      10.2.0.0/23  O_ASE   150  2         D   10.0.124.1   GigabitEthernet0/0/0

OSPF routing table status : <Inactive>
        Destinations : 0       Routes : 0
```

在R5设备上向10.2.1.1地址发送Tracert数据包。

```
<R5>tracert 10.2.1.1
 traceroute to  10.2.1.1(10.2.1.1), max hops: 30 ,packet length: 40,press CTRL_C
to break
 1 10.0.35.3 62 ms  28 ms  27 ms
 2 10.0.23.2 54 ms  58 ms  57 ms
 3  *  *  *
...
```

我们可以看到虽然Loopback 2接口被删除了，到达该目的地址的数据包仍然被R2和R3转发，直到R1上该数据包被丢弃。

步骤六. OSPF 引入缺省路由

R4的Loopback0接口连接到Internet。在R4上配置缺省路由，下一跳指向Loopback0。

```
[R4]ip route-static 0.0.0.0 0.0.0.0 LoopBack 0
```

将这条缺省路由引入到OSPF区域，定义类型为1，Cost值为10，并且定义为永久引入。

```
[R4]ospf 1
[R4-ospf-1]default-route-advertise always type 1
```

在R2上查看缺省路由的学习情况。我们可以看到R2通过第五类LSA学习到了一条默认路由，下一跳是R4的接口地址。

```
[R2]dis ip routing-table protocol ospf
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Public routing table : OSPF
         Destinations : 8       Routes : 8

OSPF routing table status : <Active>
         Destinations : 8       Routes : 8

Destination/Mask   Proto   Pre  Cost    Flags NextHop       Interface

        0.0.0.0/0  O_ASE   150  11        D  10.0.124.4   GigabitEthernet0/0/0
       10.0.1.0/24  O_ASE   150  1        D  10.0.124.1   GigabitEthernet0/0/0
       10.0.3.0/24  OSPF    10   65535    D  10.0.23.3      Serial2/0/0
       10.0.4.0/24  OSPF    10   10       D  10.0.124.4   GigabitEthernet0/0/0
       10.0.5.0/24  OSPF    10   131070   D  10.0.23.3      Serial2/0/0
      10.0.35.0/24  OSPF    10   131070   D  10.0.23.3      Serial2/0/0
       10.1.0.0/23  OSPF    10   131070   D  10.0.23.3      Serial2/0/0
       10.2.0.0/23  O_ASE   150  2        D  10.0.124.1   GigabitEthernet0/0/0

OSPF routing table status : <Inactive>
         Destinations : 0       Routes : 0
```

最后测试R5路由器与10.0.4.4之间的连通性。

```
[R5]ping -c 1 10.0.4.4
  PING 10.0.4.4: 56  data bytes, press CTRL_C to break
    Reply from 10.0.4.4: bytes=56 Sequence=1 ttl=253 time=78 ms

  --- 10.0.4.4 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 78/78/78 ms
```

步骤七. 修改 OSPF 中两类路由的优先级

查看R1的路由表，关注OSPF不同类型路由的优先级信息。

```
[R1]dis ip routing-table protocol ospf
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Public routing table : OSPF
        Destinations : 8       Routes : 8

OSPF routing table status : <Active>
        Destinations : 8       Routes : 8

Destination/Mask    Proto  Pre  Cost     Flags NextHop        Interface

       0.0.0.0/0   O_ASE  150  11        D   10.0.124.4     GigabitEthernet0/0/0
      10.0.2.0/24  OSPF   10   10        D   10.0.124.2     GigabitEthernet0/0/0
      10.0.3.0/24  OSPF   10   65545     D   10.0.124.2     GigabitEthernet0/0/0
      10.0.4.0/24  OSPF   10   10        D   10.0.124.4     GigabitEthernet0/0/0
      10.0.5.0/24  OSPF   10   131080    D   10.0.124.2     GigabitEthernet0/0/0
     10.0.23.0/24  OSPF   10   65545     D   10.0.124.2     GigabitEthernet0/0/0
     10.0.35.0/24  OSPF   10   131080    D   10.0.124.2     GigabitEthernet0/0/0
      10.1.0.0/23  OSPF   10   131080    D   10.0.124.2     GigabitEthernet0/0/0

OSPF routing table status : <Inactive>
        Destinations : 0       Routes : 0
```

默认情况下，OSPF区域内和区域之间的路由，优先级为10。OSPF外部路由，优先级为150。

修改R1和R4路由器上的OSPF区域内和区域之间的路由优先级为20，修改OSPF外部路由的优先级为50。

```
[R1]ospf
[R1-ospf-1]preference 20
[R1-ospf-1]preference ase 50

[R4]ospf
[R4-ospf-1]preference 20
[R4-ospf-1]preference ase 50
```

查看路由表中OSPF内部路由及外部路由的优先级，确认已修改成功。

```
[R1]dis ip routing-table protocol ospf
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Public routing table : OSPF
        Destinations : 8       Routes : 8

OSPF routing table status : <Active>
        Destinations : 8       Routes : 8

Destination/Mask   Proto   Pre  Cost    Flags NextHop        Interface

       0.0.0.0/0   O_ASE   50   11        D  10.0.124.4    GigabitEthernet0/0/0
      10.0.2.0/24  OSPF    20   10        D  10.0.124.2    GigabitEthernet0/0/0
      10.0.3.0/24  OSPF    20   65545     D  10.0.124.2    GigabitEthernet0/0/0
      10.0.4.0/24  OSPF    20   10        D  10.0.124.4    GigabitEthernet0/0/0
      10.0.5.0/24  OSPF    20   131080    D  10.0.124.2    GigabitEthernet0/0/0
     10.0.23.0/24  OSPF    20   65545     D  10.0.124.2    GigabitEthernet0/0/0
     10.0.35.0/24  OSPF    20   131080    D  10.0.124.2    GigabitEthernet0/0/0
      10.1.0.0/23  OSPF    20   131080    D  10.0.124.2    GigabitEthernet0/0/0

OSPF routing table status : <Inactive>
        Destinations : 0       Routes : 0
```

路由优先级仅在本地有效，用于衡量在本地通过多种方式学到的路由的优略程度。本地区域的不同路由器，如果优先级信息不同，也可以正常工作。

附加实验：思考并验证

思考在步骤六中，定义缺省路由的永久发布的作用是什么？有哪些优点和缺点？

路由汇总就像一把双刃剑，有好处也有坏处。思考并总结使用路由汇总的好处和坏处，并分析如何避免这些坏处。

最终设备配置

```
<R1>display current-configuration
[V200R001C00SPC200]
#
 sysname R1
#
interface GigabitEthernet0/0/0
 ip address 10.0.124.1 255.255.255.0
```

```
#
interface LoopBack0
 ip address 10.0.1.1 255.255.255.0
 ospf network-type broadcast
#
interface LoopBack1
 ip address 10.2.0.1 255.255.255.0
#
ospf 1 router-id 10.0.1.1
 asbr-summary 10.2.0.0 255.255.254.0
 import-route direct
 preference 20
 preference ase 50
 bandwidth-reference 10000
 area 0.0.0.2
  network 10.0.124.1 0.0.0.0
#
return

<R2>display current-configuration
[V200R001C00SPC200]
#
 sysname R2
#
interface Serial2/0/0
 link-protocol ppp
 ip address 10.0.23.2 255.255.255.0
#
interface GigabitEthernet0/0/0
 ip address 10.0.124.2 255.255.255.0
#
interface LoopBack0
 ip address 10.0.2.2 255.255.255.0
 ospf network-type broadcast
#
ospf 1 router-id 10.0.2.2
 bandwidth-reference 10000
 area 0.0.0.0
  network 10.0.23.2 0.0.0.0
  network 10.0.2.2 0.0.0.0
 area 0.0.0.2
  network 10.0.124.2 0.0.0.0
#
```

```
return

<R3>display current-configuration
[V200R001C00SPC200]
#
 sysname R3
#
interface Serial2/0/0
 link-protocol ppp
 ip address 10.0.23.3 255.255.255.0
#
interface Serial3/0/0
 link-protocol ppp
 ip address 10.0.35.3 255.255.255.0
#
interface LoopBack0
 ip address 10.0.3.3 255.255.255.0
 ospf network-type broadcast
#
ospf 1 router-id 10.0.3.3
 bandwidth-reference 10000
 area 0.0.0.0
  network 10.0.3.3 0.0.0.0
  network 10.0.23.3 0.0.0.0
 area 0.0.0.1
  abr-summary 10.1.0.0 255.255.254.0
  network 10.0.35.3 0.0.0.0
#
return

<R4>display current-configuration
[V200R001C00SPC200]
#
 sysname R4
#
interface GigabitEthernet0/0/0
 ip address 10.0.124.4 255.255.255.0
#
interface LoopBack0
 ip address 10.0.4.4 255.255.255.0
 ospf network-type broadcast
#
ospf 1 router-id 10.0.4.4
```

```
 default-route-advertise always type 1
 preference 20
 preference ase 50
 bandwidth-reference 10000
 area 0.0.0.2
  network 10.0.4.4 0.0.0.0
  network 10.0.124.4 0.0.0.0
#
 ip route-static 0.0.0.0 0.0.0.0 LoopBack0
#
return

<R5>display current-configuration
[V200R001C00SPC200]
#
 sysname R5
#
interface Serial1/0/0
 link-protocol ppp
 ip address 10.0.35.5 255.255.255.0
#
interface LoopBack0
 ip address 10.0.5.5 255.255.255.0
 ospf network-type broadcast
#
interface LoopBack1
 ip address 10.1.0.1 255.255.255.0
 ospf network-type broadcast
#
interface LoopBack2
 ip address 10.1.1.1 255.255.255.0
 ospf network-type broadcast
#
ospf 1 router-id 10.0.5.5
 bandwidth-reference 10000
 area 0.0.0.1
  network 10.0.5.5 0.0.0.0
  network 10.1.0.1 0.0.0.0
  network 10.1.1.1 0.0.0.0
  network 10.0.35.5 0.0.0.0
#
return
```

实验 2-3 OSPF 的邻接关系和 LSA

学习目的

- 了解四个OSPF邻居路由器在以太网上邻居关系建立的过程
- 掌握对OSPF的DR的选举进行干预的方法
- 观察5种类型的LSA的内容，以及它们的作用
- 了解OSPF的LSR、LSU、LSAck数据包的相互发送情况

拓扑图

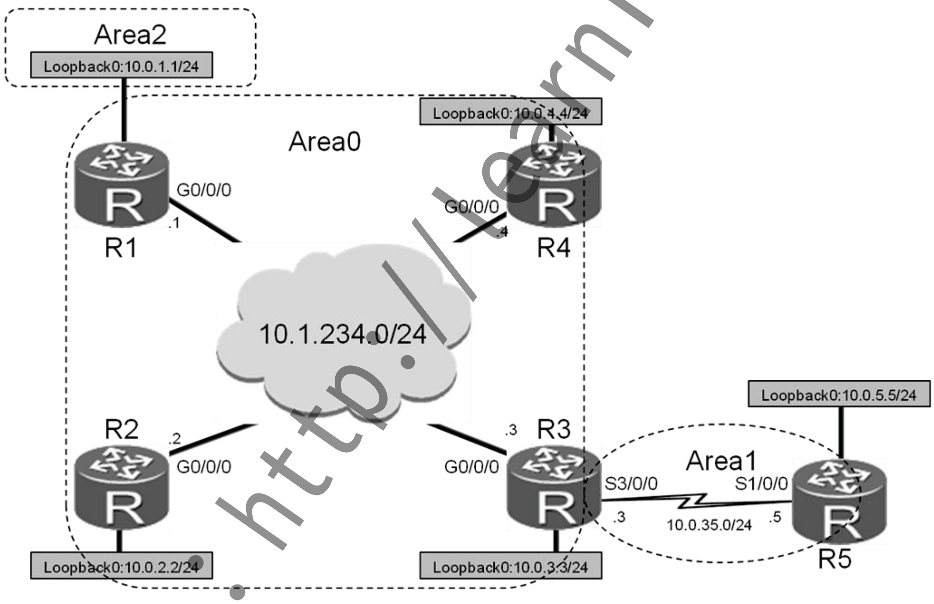

图2-3 OSPF的邻接关系和LSA

场景

你是公司的网络管理员。现在公司的网络中有五台AR G3路由器，其中R1、R2、R3和R4在公司总部，通过以太网互联。R5在公司分部，R3通过专线与公司总部的R3相连。由于网络规模较大，为了控制LSA的洪泛，你设计了多区域的OSPF互联方式。

其中R1的Loopback0接口属于区域2。R2、R3、R4的Loopback0接口与10.1.234.0/24网段属于区域0。R3与R5之间互联的网段属于区域1。R5的Loopback0接口属于OSPF外部网络。

同时为了明确设备的Router ID，你配置设备使用固定的地址作为Router ID。

在R1、R2、R3与R4之间互联的一台网络上，需要干预DR与BDR的选举。实际使用中将R3定义为DR、R2定义为BDR。R4设备定义为DROther。

学习任务

步骤一. 基础配置与 IP 编址

给所有路由器配置IP地址和掩码。配置时注意所有的Loopback接口配置掩码均为24位，模拟成一个单独的网段。

```
<R1>system-view
Enter system view, return user view with Ctrl+Z.
[R1]interface GigabitEthernet 0/0/0
[R1-GigabitEthernet0/0/0]ip address 10.1.234.1 24
[R1-GigabitEthernet0/0/0]interface LoopBack 0
[R1-LoopBack0]ip address 10.0.1.1 24

<R2>system-view
Enter system view, return user view with Ctrl+Z.
[R2]interface GigabitEthernet 0/0/0
[R2-GigabitEthernet0/0/0]ip address 10.1.234.2 24
[R2-GigabitEthernet0/0/0]interface LoopBack 0
[R2-LoopBack0]ip address 10.0.2.2 24

<R3>system-view
Enter system view, return user view with Ctrl+Z.
[R3]interface GigabitEthernet 0/0/0
[R3-GigabitEthernet0/0/0]ip address 10.1.234.3 24
[R3-GigabitEthernet0/0/0]interface Serial 3/0/0
[R3-Serial3/0/0]ip address 10.0.35.3 24
[R3-Serial3/0/0]interface LoopBack 0
[R3-LoopBack0]ip address 10.0.3.3 24

<R4>system-view
Enter system view, return user view with Ctrl+Z.
[R4]interface GigabitEthernet 0/0/0
[R4-GigabitEthernet0/0/0]ip address 10.1.234.4 24
```

```
[R4-GigabitEthernet0/0/0]interface LoopBack 0
[R4-LoopBack0]ip address 10.0.4.4 24

<R5>system-view
Enter system view, return user view with Ctrl+Z.
[R5]interface Serial 1/0/0
[R5-Serial1/0/0]ip address 10.0.35.5 24
[R5-Serial1/0/0]interface LoopBack 0
[R5-LoopBack0]ip address 10.0.5.5 24
```

配置完成后，测试直连链路的连通性。

```
[R1]ping -c 1 10.1.234.2
  PING 10.1.234.2: 56  data bytes, press CTRL_C to break
    Reply from 10.1.234.2: bytes=56 Sequence=1 ttl=255 time=13 ms

  --- 10.1.234.2 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 13/13/13 ms

[R1]ping -c 1 10.1.234.4
  PING 10.1.234.4: 56  data bytes, press CTRL_C to break
    Reply from 10.1.234.4: bytes=56 Sequence=1 ttl=255 time=6 ms

  --- 10.1.234.4 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 6/6/6 ms

[R3]ping -c 1 10.1.234.1
  PING 10.1.234.1: 56  data bytes, press CTRL_C to break
    Reply from 10.1.234.1: bytes=56 Sequence=1 ttl=255 time=13 ms

  --- 10.1.234.1 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 13/13/13 ms

[R3]ping -c 1 10.0.35.5
  PING 10.0.35.5: 56  data bytes, press CTRL_C to break
```

```
  Reply from 10.0.35.5: bytes=56 Sequence=1 ttl=255 time=32 ms

 --- 10.0.35.5 ping statistics ---
  1 packet(s) transmitted
  1 packet(s) received
  0.00% packet loss
  round-trip min/avg/max = 32/32/32 ms
```

步骤二. 配置多区域 OSPF

在R1上配置GigabitEthernet 0/0/0属于区域0，Loopback 0属于区域2。对所有OSPF区域的Loopback接口，修改其OSPF网络类型为Broadcast类型，以便于OSPF发布Loopback口的真实掩码信息。

```
[R1]ospf 1 router-id 10.0.1.1
[R1-ospf-1]area 0
[R1-ospf-1-area-0.0.0.0]network 10.1.234.1 0.0.0.0
[R1-ospf-1-area-0.0.0.0]area 2
[R1-ospf-1-area-0.0.0.2]network 10.0.1.1 0.0.0.0
[R1-ospf-1-area-0.0.0.2]interface LoopBack 0
[R1-LoopBack0]ospf network-type broadcast
```

R2、R4的所有接口均位于区域0中。

```
[R2]ospf 1 router-id 10.0.2.2
[R2-ospf-1]area 0
[R2-ospf-1-area-0.0.0.0]network 10.1.234.2 0.0.0.0
[R2-ospf-1-area-0.0.0.0]network 10.0.2.2 0.0.0.0
[R2-ospf-1-area-0.0.0.0]interface LoopBack 0
[R2-LoopBack0]ospf network-type broadcast

[R4]ospf 1 router-id 10.0.4.4
[R4-ospf-1]area 0
[R4-ospf-1-area-0.0.0.0]network 10.1.234.4 0.0.0.0
[R4-ospf-1-area-0.0.0.0]network 10.0.4.4 0.0.0.0
[R4-ospf-1-area-0.0.0.0]interface LoopBack 0
[R4-LoopBack0]ospf network-type broadcast
```

在R3上配置Loopback 0和GigabitEthernet 0/0/0属于区域0，Serial 3/0/0属于区域2。

```
[R3]ospf 1 router-id 10.0.3.3
[R3-ospf-1]area 0
```

```
[R3-ospf-1-area-0.0.0.0]network 10.1.234.3 0.0.0.0
[R3-ospf-1-area-0.0.0.0]network 10.0.3.3 0.0.0.0
[R3-ospf-1-area-0.0.0.0]area 1
[R3-ospf-1-area-0.0.0.1]network 10.0.35.3 0.0.0.0
[R3-ospf-1-area-0.0.0.1]interface LoopBack 0
[R3-LoopBack0]ospf network-type broadcast
[R3-LoopBack0]quit
```

在R5上配置Serial 1/0/0属于区域1，Loopback 0不属于任何区域。

```
[R5]osp 1 router-id 10.0.5.5
[R5-ospf-1]area 1
[R5-ospf-1-area-0.0.0.1]network 10.0.35.5 0.0.0.0
```

配置完成后，在R1查看设备的路由表。

```
[R1]display ip routing-table
Route Flags: R - relay, D - download to fib
----------------------------------------------------------------------------
Routing Tables: Public
        Destinations : 14      Routes : 14

Destination/Mask    Proto   Pre  Cost      Flags NextHop        Interface

      10.0.1.0/24   Direct  0    0           D   10.0.1.1       LoopBack0
      10.0.1.1/32   Direct  0    0           D   127.0.0.1       InLoopBack0
    10.0.1.255/32   Direct  0    0           D   127.0.0.1       InLoopBack0
      10.0.2.0/24   OSPF    10   1           D   10.1.234.2     GigabitEthernet0/0/0
      10.0.3.0/24   OSPF    10   1           D   10.1.234.3     GigabitEthernet0/0/0
      10.0.4.0/24   OSPF    10   1           D   10.1.234.4     GigabitEthernet0/0/0
     10.0.35.0/24   OSPF    10   1563        D   10.1.234.3     GigabitEthernet0/0/0
    10.1.234.0/24   Direct  0    0           D   10.1.234.1     GigabitEthernet0/0/0
    10.1.234.1/32   Direct  0    0           D   127.0.0.1       InLoopBack0
  10.1.234.255/32   Direct  0    0           D   127.0.0.1       InLoopBack0
     127.0.0.0/8    Direct  0    0           D   127.0.0.1       InLoopBack0
     127.0.0.1/32   Direct  0    0           D   127.0.0.1       InLoopBack0
127.255.255.255/32  Direct  0    0           D   127.0.0.1       InLoopBack0
255.255.255.255/32  Direct  0    0           D   127.0.0.1       InLoopBack0
```

除了没有发布进OSPF的网络10.0.5.5/24，在R1上已拥有全网的路由表。

测试网络的连通性。

```
[R1]ping -c 1 10.0.2.2
  PING 10.0.2.2: 56  data bytes, press CTRL_C to break
```

```
    Reply from 10.0.2.2: bytes=56 Sequence=1 ttl=255 time=2 ms

  --- 10.0.2.2 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 2/2/2 ms

[R1]ping -c 1 10.0.4.4
  PING 10.0.4.4: 56  data bytes, press CTRL_C to break
    Reply from 10.0.4.4: bytes=56 Sequence=1 ttl=255 time=3 ms

  --- 10.0.4.4 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 3/3/3 ms

    [R3]ping -c 1 10.0.1.1
      PING 10.0.1.1: 56  data bytes, press CTRL_C to break
        Reply from 10.0.1.1: bytes=56 Sequence=1 ttl=255 time=3 ms

      --- 10.0.1.1 ping statistics ---
        1 packet(s) transmitted
        1 packet(s) received
        0.00% packet loss
        round-trip min/avg/max = 3/3/3 ms
```

使用**display ospf brief**命令查看路由器运行的基本OSPF信息。我们可以看到，由于R1的Loopback 0接口位于区域2中，所以R1成为了一台ABR。R1的GigabitEthernet 0/0/0接口所连接的网络为广播型网络，且R1为这个网段的DR。

```
[R1]display ospf brief

         OSPF Process 1 with Router ID 10.0.1.1
                 OSPF Protocol Information

 RouterID: 10.0.1.1         Border Router:  AREA
 Multi-VPN-Instance is not enabled
 Global DS-TE Mode: Non-Standard IETF Mode
 Graceful-restart capability: disabled
 Helper support capability  : not configured
 Spf-schedule-interval: max 10000ms, start 500ms, hold 1000ms
 Default ASE parameters: Metric: 1 Tag: 1 Type: 2
```

```
Route Preference: 10
ASE Route Preference: 150
SPF Computation Count: 26
RFC 1583 Compatible
Retransmission limitation is disabled
Area Count: 2   Nssa Area Count: 0
ExChange/Loading Neighbors: 0

Area: 0.0.0.0
Authtype: None   Area flag: Normal
SPF scheduled Count: 26
ExChange/Loading Neighbors: 0

Interface: 10.1.234.1 (GigabitEthernet0/0/0)
Cost: 1       State: DR       Type: Broadcast    MTU: 1500
Priority: 1
Designated Router: 10.1.234.1
Backup Designated Router: 10.1.234.2
Timers: Hello 10 , Dead 40 , Poll  120 , Retransmit 5 , Transmit Delay 1

Area: 0.0.0.2
Authtype: None   Area flag: Normal
SPF scheduled Count: 25
ExChange/Loading Neighbors: 0

Interface: 10.0.1.1 (LoopBack0)
Cost: 0       State: DR       Type: Broadcast    MTU: 1500
Priority: 1
Designated Router: 10.0.1.1
Backup Designated Router: 0.0.0.0
Timers: Hello 10 , Dead 40 , Poll  120 , Retransmit 5 , Transmit Delay 1
```

使用**display ospf peer brief**命令查看路由器的OSPF邻居关系建立情况。由于R1是DR，所以它与该网段的所有路由器形成邻接关系。在R3上查看邻居就可以发现R3和R4之间仅存在邻居关系，而没有邻接关系。

```
[R1]display ospf peer brief

       OSPF Process 1 with Router ID 10.0.1.1
              Peer Statistic Information
----------------------------------------------------------------------------
Area Id          Interface                      Neighbor id      State
0.0.0.0          GigabitEthernet0/0/0            10.0.2.2         Full
0.0.0.0          GigabitEthernet0/0/0            10.0.3.3         Full
```

```
0.0.0.0         GigabitEthernet0/0/0          10.0.4.4        Full
----------------------------------------------------------------------------

[R3]display ospf peer brief

       OSPF Process 1 with Router ID 10.0.3.3
              Peer Statistic Information
----------------------------------------------------------------------------
Area Id         Interface                    Neighbor id      State
0.0.0.0         GigabitEthernet0/0/0          10.0.1.1        Full
0.0.0.0         GigabitEthernet0/0/0          10.0.2.2        Full
0.0.0.0         GigabitEthernet0/0/0          10.0.4.4        2-Way
0.0.0.1         Serial3/0/0                   10.0.5.5        Full
----------------------------------------------------------------------------
```

在R5上使用**display ospf lsdb**命令查看路由器的OSPF数据库信息。

```
[R5]display ospf lsdb

       OSPF Process 1 with Router ID 10.0.5.5
             Link State Database

                    Area: 0.0.0.1
Type      LinkState ID      AdvRouter          Age  Len   Sequence   Metric
Router    10.0.5.5          10.0.5.5           1182  48   80000002    1562
Router    10.0.3.3          10.0.3.3           1183  48   80000002    1562
Sum-Net   10.0.3.0          10.0.3.3           1429  28   80000001       0
Sum-Net   10.0.2.0          10.0.3.3           1429  28   80000001       1
Sum-Net   10.0.1.0          10.0.3.3           1429  28   80000001       1
Sum-Net   10.1.234.0        10.0.3.3           1429  28   80000001       1
Sum-Net   10.0.4.0          10.0.3.3           1430  28   80000001       1
```

可以看到由于在区域1中仅存在2台路由器，所以在R5的lsdb中，仅存在2条第一类LSA，剩余的5条第三类LSA是由R3向R5通告的区域间路由。

在R2上使用**display ospf lsdb**命令查看路由器的OSPF数据库信息。

```
[R2]display ospf lsdb

       OSPF Process 1 with Router ID 10.0.2.2
             Link State Database

                    Area: 0.0.0.0
Type      LinkState ID      AdvRouter          Age  Len   Sequence   Metric
Router    10.0.3.3          10.0.3.3             4  48   80000009       1
```

```
 Router    10.0.4.4         10.0.4.4          150  48     80000009      1
 Router    10.0.2.2         10.0.2.2          149  48     8000000C      1
 Router    10.0.1.1         10.0.1.1          149  36     8000000B      1
 Network   10.1.234.1       10.0.1.1          149  40     80000007      0
 Sum-Net   10.0.35.0        10.0.3.3         1790  28     80000001   1562
 Sum-Net   10.0.1.0         10.0.1.1          817  28     80000002      0
```

在R2上除了4条第一类LSA以外，还有一条第二类LSA。R2的GigabitEthernet 0/0/0所连接的是一个广播型网络，该网络上的DR会产生一条第二类LSA来描述所有的邻居。在这里可以从AdvRouter字段得知生成这条LSA的路由器是R1，符合R1是该网段DR的结果。

在R1上使用**display ospf lsdb**命令查看路由器的OSPF数据库信息。

```
[R1]display ospf lsdb

          OSPF Process 1 with Router ID 10.0.1.1
                  Link State Database

                          Area: 0.0.0.0
 Type      LinkState ID     AdvRouter          Age  Len   Sequence   Metric
 Router    10.0.3.3         10.0.3.3          447  48     80000009      1
 Router    10.0.4.4         10.0.4.4          592  48     80000009      1
 Router    10.0.2.2         10.0.2.2          592  48     8000000C      1
 Router    10.0.1.1         10.0.1.1          591  36     8000000B      1
 Network   10.1.234.1       10.0.1.1          591  40     80000007      0
 Sum-Net   10.0.35.0        10.0.3.3          434  28     80000002   1562
 Sum-Net   10.0.1.0         10.0.1.1         1259  28     80000002      0

                          Area: 0.0.0.2
 Type      LinkState ID     AdvRouter          Age  Len   Sequence   Metric
 Router    10.0.1.1         10.0.1.1         1223  36     80000004      0
 Sum-Net   10.0.35.0        10.0.1.1          433  28     80000002   1563
 Sum-Net   10.0.3.0         10.0.1.1          541  28     80000002      1
 Sum-Net   10.0.2.0         10.0.1.1          909  28     80000002      1
 Sum-Net   10.1.234.0       10.0.1.1         1269  28     80000002      1
 Sum-Net   10.0.4.0         10.0.1.1          711  28     80000002      1
```

由于R1的Loopback 0接口位于区域2中，所以R1上有2个区域的LSDB，分别是区域0和区域2的。

在R4上使用**display ospf lsdb**命令查看路由器的OSPF数据库信息。

```
[R4]display ospf lsdb
```

```
        OSPF Process 1 with Router ID 10.0.4.4
               Link State Database

                      Area: 0.0.0.0
Type      LinkState ID      AdvRouter          Age  Len    Sequence   Metric
Router    10.0.3.3          10.0.3.3           745  48     80000009       1
Router    10.0.4.4          10.0.4.4           888  48     80000009       1
Router    10.0.2.2          10.0.2.2           889  48     8000000C       1
Router    10.0.1.1          10.0.1.1           889  36     8000000B       1
Network   10.1.234.1        10.0.1.1           889  40     80000007       0
Sum-Net   10.0.35.0         10.0.3.3           732  28     80000002    1562
Sum-Net   10.0.1.0          10.0.1.1          1556  28     80000002       0
```

注意由于OSPF路由器的角色差异，OSPF链路状态数据库内容也会有所差异。对比分析R5、R2、R1与R4链路状态数据库的差异。

步骤三. 修改路由器 OSPF 接口优先级，影响 DR 选举

配置R3的G0/0/0接口优先级为255，确保R3成为10.1.234.0/24网段的DR。修改R2的G0/0/0接口优先级为254，确保R2成为10.1.234.0/24网段的BDR。修改R4的G0/0/0接口优先级为0，确保R4不参加DR/BDR选举，而成为10.1.234.0/24网段的DROther。

```
[R3]interface GigabitEthernet 0/0/0
[R3-GigabitEthernet0/0/0]ospf dr-priority 255

[R2]interface GigabitEthernet 0/0/0
[R2-GigabitEthernet0/0/0]ospf dr-priority 254

[R4]interface GigabitEthernet 0/0/0
[R4-GigabitEthernet0/0/0]ospf dr-priority 0
```

配置完成后，由于DR/BDR已经选举，并且DR/BDR角色不能抢占。所以必须关闭R1、R2、R3、R4的G0/0/0接口，并依次打开R3、R2、R1和R4的G0/0/0接口。

```
[R1-GigabitEthernet0/0/0]shutdown

[R2-GigabitEthernet0/0/0]shutdown

[R3-GigabitEthernet0/0/0]shutdown

[R4-GigabitEthernet0/0/0]shutdown
```

```
[R1-GigabitEthernet0/0/0]undo shutdown

[R2-GigabitEthernet0/0/0]undo shutdown

[R3-GigabitEthernet0/0/0]undo shutdown

[R4-GigabitEthernet0/0/0]undo shutdown
```

查看网段10.1.234.0/24网段的DR/BDR选举情况。

```
[R3]display ospf peer

        OSPF Process 1 with Router ID 10.0.3.3
               Neighbors

 Area 0.0.0.0 interface 10.1.234.3(GigabitEthernet0/0/0)'s neighbors
 Router ID: 10.0.1.1        Address: 10.1.234.1
   State: Full  Mode:Nbr is  Slave  Priority: 1
   DR: 10.1.234.3  BDR: 10.1.234.2  MTU: 0
   Dead timer due in 29  sec
   Retrans timer interval: 3
   Neighbor is up for 00:02:17
   Authentication Sequence: [ 0 ]

 Router ID: 10.0.2.2        Address: 10.1.234.2
   State: Full  Mode:Nbr is  Slave  Priority: 254
   DR: 10.1.234.3  BDR: 10.1.234.2  MTU: 0
   Dead timer due in 35  sec
   Retrans timer interval: 6
   Neighbor is up for 00:01:14
   Authentication Sequence: [ 0 ]

 Router ID: 10.0.4.4        Address: 10.1.234.4
   State: Full  Mode:Nbr is  Master  Priority: 0
   DR: 10.1.234.3  BDR: 10.1.234.2  MTU: 0
   Dead timer due in 32  sec
   Retrans timer interval: 3
   Neighbor is up for 00:01:26
   Authentication Sequence: [ 0 ]

               Neighbors

 Area 0.0.0.1 interface 10.0.35.3(Serial3/0/0)'s neighbors
```

```
 Router ID: 10.0.5.5        Address: 10.0.35.5
   State: Full  Mode:Nbr is  Master  Priority: 1
   DR: None   BDR: None   MTU: 0
   Dead timer due in 27  sec
   Retrans timer interval: 4
   Neighbor is up for 00:53:37
   Authentication Sequence: [ 0 ]
```

在重启接口后R3成为了该网段的DR，R2成为了BDR。

查看R4与R1的邻居关系。

```
[R4]display ospf peer 10.0.1.1

        OSPF Process 1 with Router ID 10.0.4.4
                Neighbors

 Area 0.0.0.0 interface 10.1.234.4(GigabitEthernet0/0/0)'s neighbors
 Router ID: 10.0.1.1        Address: 10.1.234.1
   State: 2-Way  Mode:Nbr is  Slave  Priority: 1
   DR: 10.1.234.3  BDR: 10.1.234.2  MTU: 0
   Dead timer due in 30  sec
   Retrans timer interval: 0
   Neighbor is up for 00:00:00
   Authentication Sequence: [ 0 ]
```

当邻居关系稳定以后，由于R1和R4均为DROther路由器，所以他们之间仅形成邻居关系，保持在2-way状态。

步骤四.　配置将直连路由汇总并引入到 OSPF 区域

R5的Loopback0接口不属于OSPF区域。将这条直连路由引入到OSPF区域。

```
[R5]ospf
[R5-ospf-1]import-route direct
```

在R1和R3上查看引入的外部路由。

```
[R1]display ip routing-table protocol ospf
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Public routing table : OSPF
         Destinations : 6       Routes : 6
```

```
OSPF routing table status : <Active>
        Destinations : 6       Routes : 6

Destination/Mask   Proto   Pre  Cost    Flags NextHop        Interface

      10.0.2.0/24  OSPF    10   1         D   10.1.234.2    GigabitEthernet0/0/0
      10.0.3.0/24  OSPF    10   1         D   10.1.234.3    GigabitEthernet0/0/0
      10.0.4.0/24  OSPF    10   1         D   10.1.234.4    GigabitEthernet0/0/0
      10.0.5.0/24  O_ASE   150  1         D   10.1.234.3    GigabitEthernet0/0/0
     10.0.35.0/24  OSPF    10   1563      D   10.1.234.3    GigabitEthernet0/0/0
     10.0.35.3/32  O_ASE   150  1         D   10.1.234.3    GigabitEthernet0/0/0

OSPF routing table status : <Inactive>
        Destinations : 0       Routes : 0

[R3]display ip routing-table protocol ospf
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Public routing table : OSPF
        Destinations : 5       Routes : 5

OSPF routing table status : <Active>
        Destinations : 4       Routes : 4

Destination/Mask   Proto   Pre  Cost    Flags NextHop        Interface

      10.0.1.0/24  OSPF    10   1         D   10.1.234.1    GigabitEthernet0/0/0
      10.0.2.0/24  OSPF    10   1         D   10.1.234.2    GigabitEthernet0/0/0
      10.0.4.0/24  OSPF    10   1         D   10.1.234.4    GigabitEthernet0/0/0
      10.0.5.0/24  O_ASE   150  1         D   10.0.35.5       Serial3/0/0

OSPF routing table status : <Inactive>
        Destinations : 1       Routes : 1

Destination/Mask    Proto   Pre  Cost     Flags NextHop        Interface

     10.0.35.3/32  O_ASE   150  1              10.0.35.5       Serial3/0/0
```

在R1和R3上均看到2条外部路由，分别是10.0.5.0/24和10.0.35.3/32。10.0.5.0/24为R5的Loopback接口，但为什么还有一条10.0.35.3/32呢？

查看R5的路由表，由于R3和R5之间是以PPP的形式封装的，R3的Serial 3/0/0的接口地址会以直连路由的形式出现在R5的路由表里，所以在R5上运行**import-route direct**以后该路由条目也被发布出去了。

```
[R5]display  ip routing-table
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Routing Tables: Public
        Destinations : 16       Routes : 16

Destination/Mask     Proto   Pre  Cost     Flags NextHop         Interface
     10.0.35.0/24   Direct  0    0          D   10.0.35.5       Serial1/0/0
     10.0.35.3/32   Direct  0    0          D   10.0.35.3       Serial1/0/0
     10.0.35.5/32   Direct  0    0          D   127.0.0.1       InLoopBack0
   10.0.35.255/32   Direct  0    0          D   127.0.0.1       InLoopBack0
```

最后测试网络连通性。

```
[R1]ping -c 1 10.0.5.5
  PING 10.0.5.5: 56  data bytes, press CTRL_C to break
    Reply from 10.0.5.5: bytes=56 Sequence=1 ttl=254 time=41 ms

  --- 10.0.5.5 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 41/41/41 ms
```

在R1上查看OSPF外部路由在链路状态数据库中的情况。我们可以看到，R1的LSDB中一共有3条外部路由：10.0.5.0/24、10.0.35.0/24、10.0.35.3/32。

在R1的路由表中看见的外部路由只有2条，另一条不见了。

```
 [R1]display ospf lsdb ase

        OSPF Process 1 with Router ID 10.0.1.1
                Link State Database


  Type      : External
  Ls id     : 10.0.5.0
  Adv rtr    : 10.0.5.5
  Ls age    : 834
  Len       : 36
  Options    :  E
  seq#      : 80000001
  chksum     : 0xa904
  Net mask  : 255.255.255.0
  TOS 0  Metric: 1
```

```
E type    : 2
Forwarding Address : 0.0.0.0
Tag       : 1
Priority  : Low

Type      : External
Ls id     : 10.0.35.0
Adv rtr   : 10.0.5.5
Ls age    : 1342
Len       : 36
Options   :  E
seq#      : 80000001
chksum    : 0x5e31
Net mask  : 255.255.255.0
TOS 0  Metric: 1
E type    : 2
Forwarding Address : 0.0.0.0
Tag       : 1
Priority  : Low

Type      : External
Ls id     : 10.0.35.3
Adv rtr   : 10.0.5.5
Ls age    : 1344
Len       : 36
Options   :  E
seq#      : 80000001
chksum    : 0x404c
Net mask  : 255.255.255.255
TOS 0  Metric: 1
E type    : 2
Forwarding Address : 0.0.0.0
Tag       : 1
Priority  : Medium
```

经过比较后，可以发现10.0.35.0/24这条路由是以内部路由的形式出现在路由表中的。

检查R1的LSDB中得第三类LSA就可以看到这个条目：10.0.35.0/24。

```
[R1]display ospf lsdb summary 10.0.35.0

        OSPF Process 1 with Router ID 10.0.1.1
                     Area: 0.0.0.0
```

```
             Link State Database


  Type      : Sum-Net
  Ls id     : 10.0.35.0
  Adv rtr    : 10.0.3.3
  Ls age     : 136
  Len       : 28
  Options    :  E
  seq#      : 80000004
  chksum     : 0x1ae2
  Net mask  : 255.255.255.0
  Tos 0  metric: 1562
  Priority  : Low
```

可以看出，当第三类和第五类LSA通告路由的网络位和掩码相同的情况下，OSPF优选第三类LSA通告的路由加到路由表里。

步骤五. 查看各种类型的 LSA

在R1上查看一类LSA 10.0.1.0在Area0和Area2的详细内容。

```
[R1]display ospf lsdb router 10.0.1.1

       OSPF Process 1 with Router ID 10.0.1.1
                    Area: 0.0.0.0
             Link State Database


  Type        : Router
  Ls id       : 10.0.1.1
  Adv rtr     : 10.0.1.1
  Ls age      : 591
  Len         : 36
  Options     :  ABR  E
  seq#        : 8000001e
  chksum      : 0xbc70
  Link count: 1
   * Link ID: 10.1.234.3
     Data     : 10.1.234.1
     Link Type: TransNet
     Metric : 1
                    Area: 0.0.0.2
```

```
              Link State Database


 Type       : Router
 Ls id      : 10.0.1.1
 Adv rtr    : 10.0.1.1
 Ls age     : 627
 Len        : 36
 Options    :  ABR  E
 seq#       : 80000008
 chksum     : 0x1018
 Link count: 1
  * Link ID: 10.0.1.0
    Date    : 255.255.255.0
    Link Type: StubNet
    Metric  : 0
    Priority: Low
```

对于一类LSA来说，Ls id字段表示生成这条LSA的路由器的Router ID。

R1共生成了两条第一类LSA，一条在区域0中泛洪。R1在区域0中与一个Transit网段相连，所以Link Type字段为TransNet。对于TransNet，Link ID字段为该网段上DR的接口IP地址，Data字段为本地接口的IP地址。

第二条一类LSA是R1向区域2中泛洪的，R1与区域2通过Loopback接口相连。对于Loopback接口，Link Type为StubNet，此时Link ID表示该Stub网段的IP网络地址，Data表示该Stub网段的网络掩码。

在R2、R3和R4上分别查看二类LSA 10.1.234.0在Area0的详细内容。

```
[R2]display ospf lsdb network 10.1.234.3

        OSPF Process 1 with Router ID 10.0.2.2
                     Area: 0.0.0.0
              Link State Database


 Type       : Network
 Ls id      : 10.1.234.3
 Adv rtr    : 10.0.3.3
 Ls age     : 115
 Len        : 40
 Options    :  E
 seq#       : 8000000f
 chksum     : 0x807e
```

```
  Net mask  : 255.255.255.0
  Priority  : Low
     Attached Router    10.0.3.3
     Attached Router    10.0.1.1
     Attached Router    10.0.2.2
     Attached Router    10.0.4.4
```

发现在R2、R3和R4上看到的这条LSA是一样的。

同样，可以通过Adv rtr字段得知这条LSA是由R3生成的。第二类LSA 的Ls id描述的是该网段上DR的接口IP地址，Attached Router为该网段上所有路由器的Router ID。

在R1和R3上查看三类LSA 10.0.35.0/24在Area0的详细内容。

```
[R3]display ospf lsdb summary 10.0.35.0

         OSPF Process 1 with Router ID 10.0.3.3
                          Area: 0.0.0.0
                  Link State Database


  Type      : Sum-Net
  Ls id     : 10.0.35.0
  Adv rtr   : 10.0.3.3
  Ls age    : 591
  Len       : 28
  Options   :  E
  seq#      : 8000000a
  chksum    : 0xee8
  Net mask  : 255.255.255.0
  Tos 0  metric: 1562
  Priority  : Low
```

从输出中可以看到该路由是由R3向区域0中通告的。Ls id就是通告的目的网段的网络地址，Net mask描述了目的网段的掩码信息。

```
[R1]display ospf lsdb summary 10.0.35.0

         OSPF Process 1 with Router ID 10.0.1.1
                          Area: 0.0.0.0
                  Link State Database


  Type      : Sum-Net
```

```
Ls id     : 10.0.35.0
Adv rtr    : 10.0.3.3
Ls age    : 136
Len       : 28
Options    :  E
seq#      : 80000004
chksum    : 0x1ae2
Net mask  : 255.255.255.0
Tos 0  metric: 1562
Priority  : Low
                   Area: 0.0.0.2
             Link State Database


Type      : Sum-Net
Ls id     : 10.0.35.0
Adv rtr    : 10.0.1.1
Ls age    : 382
Len       : 28
Options    :  E
seq#      : 80000002
chksum    : 0x42bf
Net mask  : 255.255.255.0
Tos 0  metric: 1563
Priority  : Low
```

在R1上共有2条描述10.0.35.0/24的第三类LSA。其中，从Adv rtr字段我们可以得知区域0中这条LSA是R3产生的。由于R1本身也是一台ABR，所以R1收到这条LSA以后又产生了一条LSA，向区域2中通告。

在R1上查看四类LSA 10.0.5.0在Area2的详细内容。第四类LSA用于描述如何到达ASBR。

```
[R1]display ospf lsdb asbr 10.0.5.5

       OSPF Process 1 with Router ID 10.0.1.1
                   Area: 0.0.0.0
             Link State Database


Type      : Sum-Asbr
Ls id     : 10.0.5.5
Adv rtr    : 10.0.3.3
Ls age    : 1119
```

```
Len        : 28
Options    :  E
seq#       : 80000008
chksum     : 0x1df3
Tos 0  metric: 1562
                   Area: 0.0.0.2
            Link State Database


Type       : Sum-Asbr
Ls id      : 10.0.5.5
Adv rtr    : 10.0.1.1
Ls age     : 1118
Len        : 28
Options    :  E
seq#       : 80000008
chksum     : 0x41d2
Tos 0  metric: 1563
```

从输出中可以看到，R1从R3收到了一条第四类LSA。Ls id用于描述ASBR的Router ID。由于这类LSA不能跨区域泛洪，所以R1又生成了一条第四类LSA向区域2中泛洪。

在R2、R4以及R3的区域0的LSDB中均存在这条LSA，因为这些路由器和ASBR（R5）不在同一个区域，他们需要通过第四类LSA来得知ASBR的位置。

```
[R2]display ospf lsdb asbr

       OSPF Process 1 with Router ID 10.0.2.2
                   Area: 0.0.0.0
            Link State Database


Type       : Sum-Asbr
Ls id      : 10.0.5.5
Adv rtr    : 10.0.3.3
Ls age     : 1676
Len        : 28
Options    :  E
seq#       : 80000008
chksum     : 0x1df3
Tos 0  metric: 1562
```

在区域1中就没有这条第四类LSA，同一个区域的路由器，不需要依赖这条

LSA来得知ASBR的位置。

步骤六. 观察 LSR、LSU 和 LSAck

我们首先观察OSPF的Update数据包及ACK数据包发送的过程。在R1上打开**debugging ospf packet update**、**debugging ospf packet ack**。

```
<R3>terminal monitor
% Current terminal monitor is on
<R3>terminal debugging
% Current terminal debugging is on
<R3>debugging ospf packet update
  OSPF Link State Update PACKET debugging switch is on
<R3>debugging ospf packet ack
  OSPF Link State Acknowledgment PACKET debugging switch is on
```

默认情况下，网络稳定运行，OSPF路由器每30分钟更新一次。为触发查询和更新信息，我们将R3的Loopback 0接口删除。

```
[R3]undo interface LoopBack 0
```

我们可以观察到首先在R1上接收到10.1.234.3发来的Update消息，消息的目的地址为224.0.0.5（即所有OSPF路由器），描述了一个网段（# Links: 1），后面有该网段的LinkID和LinkData。

```
Nov 24 2011 16:08:19.740.1+00:00 R1 RM/6/RMDEBUG:
 FileID: 0xd0178024 Line: 2178 Level: 0x20
 OSPF 1: RECV Packet. Interface: GigabitEthernet0/0/0
Nov 24 2011 16:08:19.740.2+00:00 R1 RM/6/RMDEBUG:  Source Address: 10.1.234.3
Nov 24 2011 16:08:19.740.3+00:00 R1 RM/6/RMDEBUG:  Destination Address: 224.0.0.5
Nov 24 2011 16:08:19.740.4+00:00 R1 RM/6/RMDEBUG:  Ver# 2, Type: 4 (Link-State
Update)
Nov 24 2011 16:08:19.740.5+00:00 R1 RM/6/RMDEBUG:  Length: 64, Router: 10.0.3.3
Nov 24 2011 16:08:19.740.6+00:00 R1 RM/6/RMDEBUG:  Area: 0.0.0.0, Chksum: 3671
Nov 24 2011 16:08:19.740.7+00:00 R1 RM/6/RMDEBUG:  AuType: 00
Nov 24 2011 16:08:19.740.8+00:00 R1 RM/6/RMDEBUG:  Key(ascii): 0 0 0 0 0 0 0 0
Nov 24 2011 16:08:19.740.9+00:00 R1 RM/6/RMDEBUG:  # LSAS: 1
Nov 24 2011 16:08:19.740.10+00:00 R1 RM/6/RMDEBUG:  LSA Type 1
Nov 24 2011 16:08:19.740.11+00:00 R1 RM/6/RMDEBUG:    LS ID: 10.0.3.3
Nov 24 2011 16:08:19.740.12+00:00 R1 RM/6/RMDEBUG:    Adv Rtr: 10.0.3.3
Nov 24 2011 16:08:19.740.13+00:00 R1 RM/6/RMDEBUG:    LSA Age: 1
Nov 24 2011 16:08:19.740.14+00:00 R1 RM/6/RMDEBUG:    Options: ExRouting:ON
Nov 24 2011 16:08:19.740.15+00:00 R1 RM/6/RMDEBUG:    Length: 36, Seq# 8000004f
Nov 24 2011 16:08:19.740.16+00:00 R1 RM/6/RMDEBUG:    CheckSum: 32bf
```

```
Nov 24 2011 16:08:19.740.17+00:00 R1 RM/6/RMDEBUG:    NtBit: 0  VBit: 0 EBit: 0
BBit: 1
Nov 24 2011 16:08:19.740.18+00:00 R1 RM/6/RMDEBUG:    # Links: 1
Nov 24 2011 16:08:19.740.19+00:00 R1 RM/6/RMDEBUG:    LinkID: 10.1.234.3
Nov 24 2011 16:08:19.740.20+00:00 R1 RM/6/RMDEBUG:      LinkData: 10.1.234.3
Nov 24 2011 16:08:19.740.21+00:00 R1 RM/6/RMDEBUG:      LinkType: 2
Nov 24 2011 16:08:19.740.22+00:00 R1 RM/6/RMDEBUG:      TOS# 0 Metric 1
```

接下来R1收到了一个ACK的报文，从内容看到该报文为BDR发送的ACK报文，发送的目的地址为224.0.0.5，所以R1也收到了这个报文。从该报文的序列号（Seq# 8000004f）我们可以得知，该报文为刚才Update报文的回应。

```
Nov 24 2011 16:08:20.360.1+00:00 R1 RM/6/RMDEBUG:
 FileID: 0xd0178024 Line: 2178 Level: 0x20
 OSPF 1: RECV Packet. Interface: GigabitEthernet0/0/0
Nov 24 2011 16:08:20.360.2+00:00 R1 RM/6/RMDEBUG:  Source Address: 10.1.234.2
Nov 24 2011 16:08:20.360.3+00:00 R1 RM/6/RMDEBUG:  Destination Address: 224.0.0.5
Nov 24 2011 16:08:20.360.4+00:00 R1 RM/6/RMDEBUG:  Ver# 2, Type: 5 (Link-State
Ack)
Nov 24 2011 16:08:20.360.5+00:00 R1 RM/6/RMDEBUG:  Length: 64, Router: 10.0.2.2
Nov 24 2011 16:08:20.360.6+00:00 R1 RM/6/RMDEBUG:  Area: 0.0.0.0, Chksum: 9b07
Nov 24 2011 16:08:20.360.7+00:00 R1 RM/6/RMDEBUG:  AuType: 00
Nov 24 2011 16:08:20.360.8+00:00 R1 RM/6/RMDEBUG:  Key(ascii): 0 0 0 0 0 0 0 0
Nov 24 2011 16:08:20.360.9+00:00 R1 RM/6/RMDEBUG:  # LSA Headers: 2
Nov 24 2011 16:08:20.360.10+00:00 R1 RM/6/RMDEBUG:  LSA Type 1
Nov 24 2011 16:08:20.360.11+00:00 R1 RM/6/RMDEBUG:    LS ID: 10.0.3.3
Nov 24 2011 16:08:20.360.12+00:00 R1 RM/6/RMDEBUG:    Adv Rtr: 10.0.3.3
Nov 24 2011 16:08:20.360.13+00:00 R1 RM/6/RMDEBUG:    LSA Age: 6
Nov 24 2011 16:08:20.360.14+00:00 R1 RM/6/RMDEBUG:    Options: ExRouting:ON
Nov 24 2011 16:08:20.360.15+00:00 R1 RM/6/RMDEBUG:    Length: 36, Seq# 8000004f
Nov 24 2011 16:08:20.360.16+00:00 R1 RM/6/RMDEBUG:    CheckSum: 32bf
```

最后是R1自己发送的ACK报文。报文源地址为R1 GigabitEthernet 0/0/0的接口地址，目的地址为224.0.0.6。该报文是发送给DR和BDR的。该报文的序列号也是8000004f。

```
Nov 24 2011 16:08:20.650.1+00:00 R1 RM/6/RMDEBUG:
 FileID: 0xd0178025 Line: 4383 Level: 0x20
 OSPF 1: SEND Packet. Interface: GigabitEthernet0/0/0
Nov 24 2011 16:08:20.650.2+00:00 R1 RM/6/RMDEBUG:  Source Address: 10.1.234.1
Nov 24 2011 16:08:20.650.3+00:00 R1 RM/6/RMDEBUG:  Destination Address: 224.0.0.6
Nov 24 2011 16:08:20.650.4+00:00 R1 RM/6/RMDEBUG:  Ver# 2, Type: 5 (Link-State
Ack)
Nov 24 2011 16:08:20.650.5+00:00 R1 RM/6/RMDEBUG:  Length: 44, Router: 10.0.1.1
```

```
Nov 24 2011 16:08:20.650.6+00:00 R1 RM/6/RMDEBUG:  Area: 0.0.0.0, Chksum: 2392
Nov 24 2011 16:08:20.650.7+00:00 R1 RM/6/RMDEBUG:  AuType: 00
Nov 24 2011 16:08:20.650.8+00:00 R1 RM/6/RMDEBUG:  Key(ascii): 0 0 0 0 0 0 0 0
Nov 24 2011 16:08:20.650.9+00:00 R1 RM/6/RMDEBUG:  # LSA Headers: 1
Nov 24 2011 16:08:20.650.10+00:00 R1 RM/6/RMDEBUG:  LSA Type 1
Nov 24 2011 16:08:20.650.11+00:00 R1 RM/6/RMDEBUG:    LS ID: 10.0.3.3
Nov 24 2011 16:08:20.650.12+00:00 R1 RM/6/RMDEBUG:    Adv Rtr: 10.0.3.3
Nov 24 2011 16:08:20.650.13+00:00 R1 RM/6/RMDEBUG:    LSA Age: 2
Nov 24 2011 16:08:20.650.14+00:00 R1 RM/6/RMDEBUG:    Options: ExRouting:ON
Nov 24 2011 16:08:20.650.15+00:00 R1 RM/6/RMDEBUG:    Length: 36, Seq# 8000004f
Nov 24 2011 16:08:20.650.16+00:00 R1 RM/6/RMDEBUG:    CheckSum: 32bf
```

接下来恢复R3上删除的Loopback0接口。

```
[R3]interface loopback 0
[R3-LoopBack0]ip address 10.0.3.3 24
```

和刚才一样，R1首先收到来自R3的Update报文，但这次在报文中通告了一个新的网段，所以# Links这里值为2，后面有新通告的网段的网络号和掩码。

```
<R1>
Nov 24 2011 16:13:50.110.1+00:00 R1 RM/6/RMDEBUG:
 FileID: 0xd0178024 Line: 2178 Level: 0x20
 OSPF 1: RECV Packet. Interface: GigabitEthernet0/0/0
Nov 24 2011 16:13:50.110.2+00:00 R1 RM/6/RMDEBUG:  Source Address: 10.1.234.3
Nov 24 2011 16:13:50.110.3+00:00 R1 RM/6/RMDEBUG:  Destination Address: 224.0.0.5
Nov 24 2011 16:13:50.110.4+00:00 R1 RM/6/RMDEBUG:  Ver# 2, Type: 4 (Link-State
Update)
Nov 24 2011 16:13:50.110.5+00:00 R1 RM/6/RMDEBUG:  Length: 76, Router: 10.0.3.3
Nov 24 2011 16:13:50.110.6+00:00 R1 RM/6/RMDEBUG:  Area: 0.0.0.0, Chksum: 8516
Nov 24 2011 16:13:50.110.7+00:00 R1 RM/6/RMDEBUG:  AuType: 00
Nov 24 2011 16:13:50.110.8+00:00 R1 RM/6/RMDEBUG:  Key(ascii): 0 0 0 0 0 0 0 0
Nov 24 2011 16:13:50.110.9+00:00 R1 RM/6/RMDEBUG:  # LSAS: 1
Nov 24 2011 16:13:50.110.10+00:00 R1 RM/6/RMDEBUG:  LSA Type 1
Nov 24 2011 16:13:50.110.11+00:00 R1 RM/6/RMDEBUG:    LS ID: 10.0.3.3
Nov 24 2011 16:13:50.110.12+00:00 R1 RM/6/RMDEBUG:    Adv Rtr: 10.0.3.3
Nov 24 2011 16:13:50.110.13+00:00 R1 RM/6/RMDEBUG:    LSA Age: 1
Nov 24 2011 16:13:50.110.14+00:00 R1 RM/6/RMDEBUG:    Options: ExRouting:ON
Nov 24 2011 16:13:50.110.15+00:00 R1 RM/6/RMDEBUG:    Length: 48, Seq# 80000056
Nov 24 2011 16:13:50.110.16+00:00 R1 RM/6/RMDEBUG:    CheckSum: d3f6
Nov 24 2011 16:13:50.110.17+00:00 R1 RM/6/RMDEBUG:    NtBit: 0  VBit: 0 EBit: 0
BBit: 1
Nov 24 2011 16:13:50.110.18+00:00 R1 RM/6/RMDEBUG:    # Links: 2
Nov 24 2011 16:13:50.110.19+00:00 R1 RM/6/RMDEBUG:    LinkID: 10.1.234.3
```

```
Nov 24 2011 16:13:50.110.20+00:00 R1 RM/6/RMDEBUG:     LinkData: 10.1.234.3
Nov 24 2011 16:13:50.110.21+00:00 R1 RM/6/RMDEBUG:     LinkType: 2
Nov 24 2011 16:13:50.110.22+00:00 R1 RM/6/RMDEBUG:     TOS# 0 Metric 1
Nov 24 2011 16:13:50.110.23+00:00 R1 RM/6/RMDEBUG:   LinkID: 10.0.3.3
Nov 24 2011 16:13:50.110.24+00:00 R1 RM/6/RMDEBUG:   LinkData: 255.255.255.255
Nov 24 2011 16:13:50.110.25+00:00 R1 RM/6/RMDEBUG:     LinkType: 3
Nov 24 2011 16:13:50.110.26+00:00 R1 RM/6/RMDEBUG:     TOS# 0 Metric 0
Nov 24 2011 16:13:50.110.27+00:00 R1 RM/6/RMDEBUG:
Nov 24 2011 16:13:50.360.1+00:00 R1 RM/6/RMDEBUG:
```

R1首先收到BDR的ACK报文。

```
 FileID: 0xd0178024 Line: 2178 Level: 0x20
 OSPF 1: RECV Packet. Interface: GigabitEthernet0/0/0
Nov 24 2011 16:13:50.360.2+00:00 R1 RM/6/RMDEBUG:  Source Address: 10.1.234.2
Nov 24 2011 16:13:50.360.3+00:00 R1 RM/6/RMDEBUG:  Destination Address: 224.0.0.5
Nov 24 2011 16:13:50.360.4+00:00 R1 RM/6/RMDEBUG:  Ver# 2, Type: 5 (Link-State
Ack)
Nov 24 2011 16:13:50.360.5+00:00 R1 RM/6/RMDEBUG:  Length: 44, Router: 10.0.2.2
Nov 24 2011 16:13:50.360.6+00:00 R1 RM/6/RMDEBUG:  Area: 0.0.0.0, Chksum: 8147
Nov 24 2011 16:13:50.360.7+00:00 R1 RM/6/RMDEBUG:  AuType: 00
Nov 24 2011 16:13:50.360.8+00:00 R1 RM/6/RMDEBUG:  Key(ascii): 0 0 0 0 0 0 0 0
Nov 24 2011 16:13:50.360.9+00:00 R1 RM/6/RMDEBUG:  # LSA Headers: 1
Nov 24 2011 16:13:50.360.10+00:00 R1 RM/6/RMDEBUG:  LSA Type 1
Nov 24 2011 16:13:50.360.11+00:00 R1 RM/6/RMDEBUG:    LS ID: 10.0.3.3
Nov 24 2011 16:13:50.360.12+00:00 R1 RM/6/RMDEBUG:    Adv Rtr: 10.0.3.3
Nov 24 2011 16:13:50.360.13+00:00 R1 RM/6/RMDEBUG:    LSA Age: 1
Nov 24 2011 16:13:50.360.14+00:00 R1 RM/6/RMDEBUG:    Options: ExRouting:ON
Nov 24 2011 16:13:50.360.15+00:00 R1 RM/6/RMDEBUG:    Length: 48, Seq# 80000056
Nov 24 2011 16:13:50.360.16+00:00 R1 RM/6/RMDEBUG:    CheckSum: d3f6
Nov 24 2011 16:13:50.360.17+00:00 R1 RM/6/RMDEBUG:
Nov 24 2011 16:13:50.570.1+00:00 R1 RM/6/RMDEBUG:
```

最后是R1自己发送的ACK报文。

```
 FileID: 0xd0178025 Line: 4383 Level: 0x20
 OSPF 1: SEND Packet. Interface: GigabitEthernet0/0/0
Nov 24 2011 16:13:50.570.2+00:00 R1 RM/6/RMDEBUG:  Source Address: 10.1.234.1
Nov 24 2011 16:13:50.570.3+00:00 R1 RM/6/RMDEBUG:  Destination Address: 224.0.0.6
Nov 24 2011 16:13:50.570.4+00:00 R1 RM/6/RMDEBUG:  Ver# 2, Type: 5 (Link-State
Ack)
Nov 24 2011 16:13:50.570.5+00:00 R1 RM/6/RMDEBUG:  Length: 44, Router: 10.0.1.1
Nov 24 2011 16:13:50.570.6+00:00 R1 RM/6/RMDEBUG:  Area: 0.0.0.0, Chksum: 8248
Nov 24 2011 16:13:50.570.7+00:00 R1 RM/6/RMDEBUG:  AuType: 00
```

```
Nov 24 2011 16:13:50.570.8+00:00 R1 RM/6/RMDEBUG:  Key(ascii): 0 0 0 0 0 0 0 0
Nov 24 2011 16:13:50.570.9+00:00 R1 RM/6/RMDEBUG:  # LSA Headers: 1
Nov 24 2011 16:13:50.570.10+00:00 R1 RM/6/RMDEBUG:  LSA Type 1
Nov 24 2011 16:13:50.570.11+00:00 R1 RM/6/RMDEBUG:    LS ID: 10.0.3.3
Nov 24 2011 16:13:50.570.12+00:00 R1 RM/6/RMDEBUG:    Adv Rtr: 10.0.3.3
Nov 24 2011 16:13:50.570.13+00:00 R1 RM/6/RMDEBUG:    LSA Age: 1
Nov 24 2011 16:13:50.570.14+00:00 R1 RM/6/RMDEBUG:    Options: ExRouting:ON
Nov 24 2011 16:13:50.570.15+00:00 R1 RM/6/RMDEBUG:    Length: 48, Seq# 80000056
Nov 24 2011 16:13:50.570.16+00:00 R1 RM/6/RMDEBUG:    CheckSum: d3f6
```

在下面一个步骤中我们看Request报文。正常情况下，路由器不会主动发送该报文，为观察该报文的发送，我们将R1的OSPF进程重启。首先在路由器上观察到的是R1向R2发起了LS Request。

```
<R1>reset ospf process
Warning: The OSPF process will be reset. Continue? [Y/N]:y
Nov 24 2011 16:31:42.270.1+00:00 R1 RM/6/RMDEBUG:
 FileID: 0xd0178025 Line: 2842 Level: 0x20
 OSPF 1: SEND Packet. Interface: GigabitEthernet0/0/0
Nov 24 2011 16:31:42.270.2+00:00 R1 RM/6/RMDEBUG:  Source Address: 10.1.234.1
Nov 24 2011 16:31:42.270.3+00:00 R1 RM/6/RMDEBUG:  Destination Address:
10.1.234.2
Nov 24 2011 16:31:42.270.4+00:00 R1 RM/6/RMDEBUG:  Ver# 2, Type: 3 (Link-State
Req)
Nov 24 2011 16:31:42.270.5+00:00 R1 RM/6/RMDEBUG:  Length: 144, Router: 10.0.1.1
Nov 24 2011 16:31:42.270.6+00:00 R1 RM/6/RMDEBUG:  Area: 0.0.0.0, Chksum: a316
Nov 24 2011 16:31:42.270.7+00:00 R1 RM/6/RMDEBUG:  AuType: 00
Nov 24 2011 16:31:42.270.8+00:00 R1 RM/6/RMDEBUG:  Key(ascii): 0 0 0 0 0 0 0 0
Nov 24 2011 16:31:42.270.9+00:00 R1 RM/6/RMDEBUG:  # Requesting LSAs: 10
Nov 24 2011 16:31:42.270.10+00:00 R1 RM/6/RMDEBUG:  LSA Type 1
Nov 24 2011 16:31:42.270.11+00:00 R1 RM/6/RMDEBUG:    LS ID: 10.0.1.1
Nov 24 2011 16:31:42.270.12+00:00 R1 RM/6/RMDEBUG:    Adv Rtr: 10.0.1.1
Nov 24 2011 16:31:42.270.13+00:00 R1 RM/6/RMDEBUG:  LSA Type 1
Nov 24 2011 16:31:42.270.14+00:00 R1 RM/6/RMDEBUG:    LS ID: 10.0.3.3
Nov 24 2011 16:31:42.270.15+00:00 R1 RM/6/RMDEBUG:    Adv Rtr: 10.0.3.3
Nov 24 2011 16:31:42.280.1+00:00 R1 RM/6/RMDEBUG:  LSA Type 1
Nov 24 2011 16:31:42.280.2+00:00 R1 RM/6/RMDEBUG:    LS ID: 10.0.4.4
Nov 24 2011 16:31:42.280.3+00:00 R1 RM/6/RMDEBUG:    Adv Rtr: 10.0.4.4
Nov 24 2011 16:31:42.280.4+00:00 R1 RM/6/RMDEBUG:  LSA Type 2
Nov 24 2011 16:31:42.280.5+00:00 R1 RM/6/RMDEBUG:    LS ID: 10.1.234.3
Nov 24 2011 16:31:42.280.6+00:00 R1 RM/6/RMDEBUG:    Adv Rtr: 10.0.3.3
Nov 24 2011 16:31:42.280.7+00:00 R1 RM/6/RMDEBUG:  LSA Type 3
Nov 24 2011 16:31:42.280.8+00:00 R1 RM/6/RMDEBUG:    LS ID: 10.0.35.0
Nov 24 2011 16:31:42.280.9+00:00 R1 RM/6/RMDEBUG:    Adv Rtr: 10.0.3.3
```

```
Nov 24 2011 16:31:42.280.10+00:00 R1 RM/6/RMDEBUG:  LSA Type 3
Nov 24 2011 16:31:42.280.11+00:00 R1 RM/6/RMDEBUG:    LS ID: 10.0.1.0
Nov 24 2011 16:31:42.280.12+00:00 R1 RM/6/RMDEBUG:    Adv Rtr: 10.0.1.1
Nov 24 2011 16:31:42.280.13+00:00 R1 RM/6/RMDEBUG:  LSA Type 4
Nov 24 2011 16:31:42.280.14+00:00 R1 RM/6/RMDEBUG:    LS ID: 10.0.5.5
Nov 24 2011 16:31:42.280.15+00:00 R1 RM/6/RMDEBUG:    Adv Rtr: 10.0.3.3
Nov 24 2011 16:31:42.280.16+00:00 R1 RM/6/RMDEBUG:  LSA Type 5
Nov 24 2011 16:31:42.280.17+00:00 R1 RM/6/RMDEBUG:    LS ID: 10.0.35.0
Nov 24 2011 16:31:42.280.18+00:00 R1 RM/6/RMDEBUG:    Adv Rtr: 10.0.5.5
Nov 24 2011 16:31:42.280.19+00:00 R1 RM/6/RMDEBUG:  LSA Type 5
Nov 24 2011 16:31:42.280.20+00:00 R1 RM/6/RMDEBUG:    LS ID: 10.0.35.3
Nov 24 2011 16:31:42.280.21+00:00 R1 RM/6/RMDEBUG:    Adv Rtr: 10.0.5.5
Nov 24 2011 16:31:42.280.22+00:00 R1 RM/6/RMDEBUG:  LSA Type 5
Nov 24 2011 16:31:42.280.23+00:00 R1 RM/6/RMDEBUG:    LS ID: 10.0.5.0
Nov 24 2011 16:31:42.280.24+00:00 R1 RM/6/RMDEBUG:    Adv Rtr: 10.0.5.5
```

随后R1收到了来自R3的LS Request。

```
Nov 24 2011 16:31:48.320.1+00:00 R1 RM/6/RMDEBUG:
 FileID: 0xd0178024 Line: 2178 Level: 0x20
 OSPF 1: RECV Packet. Interface: GigabitEthernet0/0/0
Nov 24 2011 16:31:48.320.2+00:00 R1 RM/6/RMDEBUG:  Source Address: 10.1.234.3
Nov 24 2011 16:31:48.320.3+00:00 R1 RM/6/RMDEBUG:  Destination Address:
10.1.234.1
Nov 24 2011 16:31:48.320.4+00:00 R1 RM/6/RMDEBUG:  Ver# 2, Type: 3 (Link-State
Req)
Nov 24 2011 16:31:48.320.5+00:00 R1 RM/6/RMDEBUG:  Length: 48, Router: 10.0.3.3
Nov 24 2011 16:31:48.320.6+00:00 R1 RM/6/RMDEBUG:  Area: 0.0.0.0, Chksum: c4c2
Nov 24 2011 16:31:48.320.7+00:00 R1 RM/6/RMDEBUG:  AuType: 00
Nov 24 2011 16:31:48.320.8+00:00 R1 RM/6/RMDEBUG:  Key(ascii): 0 0 0 0 0 0 0 0
Nov 24 2011 16:31:48.320.9+00:00 R1 RM/6/RMDEBUG:  # Requesting LSAs: 2
Nov 24 2011 16:31:48.320.10+00:00 R1 RM/6/RMDEBUG:  LSA Type 1
Nov 24 2011 16:31:48.320.11+00:00 R1 RM/6/RMDEBUG:    LS ID: 10.0.1.1
Nov 24 2011 16:31:48.320.12+00:00 R1 RM/6/RMDEBUG:    Adv Rtr: 10.0.1.1
Nov 24 2011 16:31:48.320.13+00:00 R1 RM/6/RMDEBUG:  LSA Type 3
Nov 24 2011 16:31:48.320.14+00:00 R1 RM/6/RMDEBUG:    LS ID: 10.0.1.0
Nov 24 2011 16:31:48.320.15+00:00 R1 RM/6/RMDEBUG:    Adv Rtr: 10.0.1.1
```

附加实验：思考并验证

假设区域2存在一台路由器R6。它计算到达10.0.5.0/24网段的路由信息与R2、R3计算该信息的步骤有什么差异？

类型4的LSA什么时候会出现？

实验中如果将R1和R4都配置成DROther，会有什么隐患？

最终设备配置

```
<R1>display current-configuration
[V200R001C00SPC200]
#
 sysname R1
#
interface GigabitEthernet0/0/0
 ip address 10.1.234.1 255.255.255.0
#
interface LoopBack0
 ip address 10.0.1.1 255.255.255.0
 ospf network-type broadcast
#
ospf 1 router-id 10.0.1.1
 area 0.0.0.0
  network 10.1.234.1 0.0.0.0
 area 0.0.0.2
  network 10.0.1.1 0.0.0.0
#
return

<R2>display current-configuration
[V200R001C00SPC200]
#
 sysname R2
#
interface GigabitEthernet0/0/0
 ip address 10.1.234.2 255.255.255.0
 ospf dr-priority 254
#
interface LoopBack0
 ip address 10.0.2.2 255.255.255.0
 ospf network-type broadcast
#
ospf 1 router-id 10.0.2.2
 area 0.0.0.0
  network 10.1.234.2 0.0.0.0
  network 10.0.2.2 0.0.0.0
```

```
#
return

<R3>display current-configuration
[V200R001C00SPC200]
#
 sysname R3
#
interface Serial3/0/0
 link-protocol ppp
 ip address 10.0.35.3 255.255.255.0
#
interface GigabitEthernet0/0/0
 ip address 10.1.234.3 255.255.255.0
 ospf dr-priority 255
#
interface LoopBack0
 ip address 10.0.3.3 255.255.255.0
 ospf network-type broadcast
#
ospf 1 router-id 10.0.3.3
 area 0.0.0.0
  network 10.1.234.3 0.0.0.0
  network 10.0.3.3 0.0.0.0
 area 0.0.0.1
  network 10.0.35.3 0.0.0.0
#
return

<R4>display current-configuration
[V200R001C00SPC200]
#
 sysname R4
#
interface GigabitEthernet0/0/0
 ip address 10.1.234.4 255.255.255.0
 ospf dr-priority 0
#
interface LoopBack0
 ip address 10.0.4.4 255.255.255.0
 ospf network-type broadcast
#
ospf 1 router-id 10.0.4.4
```

```
 area 0.0.0.0
  network 10.1.234.4 0.0.0.0
  network 10.0.4.4 0.0.0.0
#
return

<R5>display current-configuration
[V200R001C00SPC200]
#
 sysname R5
#
interface Serial1/0/0
 link-protocol ppp
 ip address 10.0.35.5 255.255.255.0
#
interface LoopBack0
 ip address 10.0.5.5 255.255.255.0
#
ospf 1 router-id 10.0.5.5
 import-route direct
 area 0.0.0.1
  network 10.0.35.5 0.0.0.0
#
return
```

实验 2-4 OSPF Stub 区域与 NSSA 区域

学习目的

- 掌握OSPF的Stub区域的配置
- 掌握OSPF的NSSA区域的配置
- 观察LSA Type7的内容
- 理解LSA Type7与Type5之间的转化关系

拓扑图

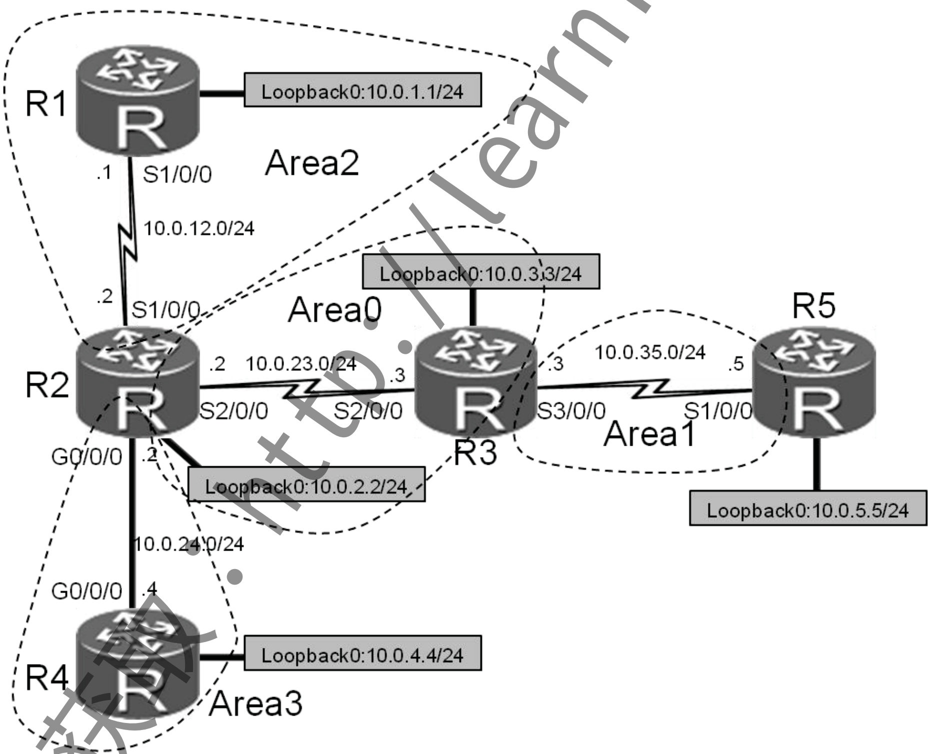

图2-4 OSPF Stub区域与NSSA区域

场景

你是公司的网络管理员。现在公司的网络中有五台AR G3路由器，其中R2、R3和R4在公司总部。R5在公司分部。R3通过专线与公司总部的R3相连。R1在公司的另外一个分部，通过专线与公司总部的R2相连。

网段10.0.23.0/24、10.0.2.0/24、10.0.3.0/24属于区域0。

网段10.0.35.0/24属于区域1，区域1为NSSA区域。R5的Loopback0接口不属于OSPF区域。

网段10.0.24.0/24属于区域3。R4的Loopback0接口连接到Internet，需要配置一条缺省路由。

网段10.0.12.0/24、10.0.1.0/24属于区域2，区域2为Stub区域。

同时为了明确设备的Router-ID，你配置设备使用固定的地址作为Router-ID。

学习任务

步骤一. 基础配置与 IP 编址

给所有路由器配置IP地址和掩码。配置时注意所有的Loopback接口配置掩码均为24位，模拟成一个单独的网段。

```
<R1>system-view
Enter system view, return user view with Ctrl+Z.
[R1]interface Serial 1/0/0
[R1-Serial1/0/0]ip address 10.0.12.1 24
[R1-Serial1/0/0]quit
[R1]interface LoopBack 0
[R1-LoopBack0]ip address 10.0.1.1 24
[R1-LoopBack0]quit

<R2>system-view
Enter system view, return user view with Ctrl+Z.
[R2]interface Serial 1/0/0
[R2-Serial1/0/0]ip address 10.0.12.2 24
[R2-Serial1/0/0]quit
[R2]interface Serial 2/0/0
[R2-Serial2/0/0]ip address 10.0.23.2 24
[R2-Serial2/0/0]quit
[R2]interface GigabitEthernet 0/0/0
```

```
[R2-GigabitEthernet0/0/0]ip address 10.0.24.2 24
[R2-GigabitEthernet0/0/0]quit
[R2]interface LoopBack 0
[R2-LoopBack0]ip address 10.0.2.2 24
[R2-LoopBack0]quit

<R3>system-view
Enter system view, return user view with Ctrl+Z.
[R3]interface Serial 2/0/0
[R3-Serial2/0/0]ip address 10.0.23.3 24
[R3-Serial2/0/0]quit
[R3]interface Serial 3/0/0
[R3-Serial3/0/0]ip address 10.0.35.3 24
[R3-Serial3/0/0]quit
[R3]interface LoopBack 0
[R3-LoopBack0]ip address 10.0.3.3 24
[R3-LoopBack0]quit

<R4>system-view
Enter system view, return user view with Ctrl+Z.
[R4]interface GigabitEthernet 0/0/0
[R4-GigabitEthernet0/0/0]ip address 10.0.24.4 24
[R4-GigabitEthernet0/0/0]quit
[R4]interface LoopBack 0
[R4-LoopBack0]ip address 10.0.4.4 24
[R4-LoopBack0]quit

<R5>system-view
Enter system view, return user view with Ctrl+Z.
[R5]interface Serial 1/0/0
[R5-Serial1/0/0]ip address 10.0.35.5 24
[R5-Serial1/0/0]quit
[R5]interface LoopBack 0
[R5-LoopBack0]ip address 10.0.5.5 24
[R5-LoopBack0]quit
```

测试直连链路的连通性。

```
[R2]ping -c 1 10.0.12.1
  PING 10.0.12.1: 56  data bytes, press CTRL_C to break
    Reply from 10.0.12.1: bytes=56 Sequence=1 ttl=255 time=30 ms

  --- 10.0.12.1 ping statistics ---
    1 packet(s) transmitted
```

```
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 30/30/30 ms

[R2]ping -c 1 10.0.24.4
  PING 10.0.24.4: 56  data bytes, press CTRL_C to break
    Reply from 10.0.24.4: bytes=56 Sequence=1 ttl=255 time=6 ms

  --- 10.0.24.4 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 6/6/6 ms

[R2]ping -c 1 10.0.23.3
  PING 10.0.23.3: 56  data bytes, press CTRL_C to break
    Reply from 10.0.23.3: bytes=56 Sequence=1 ttl=255 time=31 ms

  --- 10.0.23.3 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 31/31/31 ms

[R3]ping -c 1 10.0.35.5
  PING 10.0.35.5: 56  data bytes, press CTRL_C to break
    Reply from 10.0.35.5: bytes=56 Sequence=1 ttl=255 time=38 ms

  --- 10.0.35.5 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 38/38/38 ms
```

步骤二. 配置多区域 OSPF

在R1上配置Serial 1/0/0及Loopback 0属于区域0，并对所有OSPF区域的Loopback接口，修改其OSPF网络类型为Broadcast类型，以便于OSPF发布Loopback口的真实掩码信息，所有的路由器使用Loopback 0的IP地址作为Router ID。

```
[R1]ospf 1 router-id 10.0.1.1
```

```
[R1-ospf-1]area 2
[R1-ospf-1-area-0.0.0.2]network 10.0.12.1 0.0.0.0
[R1-ospf-1-area-0.0.0.2]network 10.0.1.1 0.0.0.0
[R1-ospf-1-area-0.0.0.2]quit
[R1-ospf-1]quit
[R1]interface LoopBack 0
[R1-LoopBack0]ospf network-type broadcast
[R1-LoopBack0]quit
```

在R2上配置接口Serial 2/0/0及Loopback 0属于区域0，接口Serial 1/0/0属于区域2，接口GigabitEthernet 0/0/0属于区域3。

```
[R2]ospf 1 router-id 10.0.2.2
[R2-ospf-1]area 0
[R2-ospf-1-area-0.0.0.0]network 10.0.23.2 0.0.0.0
[R2-ospf-1-area-0.0.0.0]network 10.0.2.2 0.0.0.0
[R2-ospf-1-area-0.0.0.0]quit
[R2-ospf-1]area 2
[R2-ospf-1-area-0.0.0.2]network 10.0.12.2 0.0.0.0
[R2-ospf-1-area-0.0.0.2]quit
[R2-ospf-1]area 3
[R2-ospf-1-area-0.0.0.3]network 10.0.24.2 0.0.0.0
[R2-ospf-1-area-0.0.0.3]qui
[R2-ospf-1]quit
[R2]int LoopBack 0
[R2-LoopBack0]ospf network-type broadcast
[R2-LoopBack0]quit
```

在R3上配置接口Serial 2/0/0及Loopback 0属于区域0，接口Serial 3/0/0属于区域1。

```
[R3]ospf 1 router-id 10.0.3.3
[R3-ospf-1]area 0
[R3-ospf-1-area-0.0.0.0]network 10.0.23.3 0.0.0.0
[R3-ospf-1-area-0.0.0.0]network 10.0.3.3 0.0.0.0
[R3-ospf-1-area-0.0.0.0]quit
[R3-ospf-1]area 1
[R3-ospf-1-area-0.0.0.1]network 10.0.35.3 0.0.0.0
[R3-ospf-1-area-0.0.0.1]quit
[R3-ospf-1]quit
[R3]interface LoopBack 0
[R3-LoopBack0]ospf network-type broadcast
[R3-LoopBack0]quit
```

在R4上配置接口GigabitEthernet 0/0/0属于区域3，接口Loopback 0不属于任何区域。

```
[R4]ospf 1 router-id 10.0.4.4
[R4-ospf-1]area 3
[R4-ospf-1-area-0.0.0.3]network 10.0.24.4  0.0.0.0
[R4-ospf-1-area-0.0.0.3]quit
[R4-ospf-1]quit
```

在R5上配置接口Serial 1/0/0属于区域1，接口Loopback 0不属于OSPF区域。

```
[R5]ospf 1 router-id 10.0.5.5
[R5-ospf-1]area 1
[R5-ospf-1-area-0.0.0.1]network 10.0.35.5 0.0.0.0
[R5-ospf-1-area-0.0.0.1]quit
[R5-ospf-1]quit
```

配置完成后，在R1上查看路由器的路由表。

```
[R1]display ip routing-table
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Routing Tables: Public
        Destinations : 16       Routes : 16

Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface

      10.0.1.0/24   Direct  0    0           D   10.0.1.1        LoopBack0
      10.0.1.1/32   Direct  0    0           D   127.0.0.1       InLoopBack0
    10.0.1.255/32   Direct  0    0           D   127.0.0.1       InLoopBack0
      10.0.2.0/24   OSPF    10   1562        D   10.0.12.2       Serial1/0/0
      10.0.3.0/24   OSPF    10   3124        D   10.0.12.2       Serial1/0/0
     10.0.12.0/24   Direct  0    0           D   10.0.12.1       Serial1/0/0
     10.0.12.1/32   Direct  0    0           D   127.0.0.1       InLoopBack0
     10.0.12.2/32   Direct  0    0           D   10.0.12.2       Serial1/0/0
   10.0.12.255/32   Direct  0    0           D   127.0.0.1       InLoopBack0
     10.0.23.0/24   OSPF    10   3124        D   10.0.12.2       Serial1/0/0
     10.0.24.0/24   OSPF    10   1563        D   10.0.12.2       Serial1/0/0
     10.0.35.0/24   OSPF    10   4686        D   10.0.12.2       Serial1/0/0
      127.0.0.0/8   Direct  0    0           D   127.0.0.1       InLoopBack0
      127.0.0.1/32  Direct  0    0           D   127.0.0.1       InLoopBack0
127.255.255.255/32  Direct  0    0           D   127.0.0.1       InLoopBack0
255.255.255.255/32  Direct  0    0           D   127.0.0.1       InLoopBack0
```

测试全网的连通性。

```
[R1]ping -c 1 10.0.35.5
  PING 10.0.35.5: 56  data bytes, press CTRL_C to break
    Reply from 10.0.35.5: bytes=56 Sequence=1 ttl=253 time=114 ms

  --- 10.0.35.5 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 114/114/114 ms

[R1]ping -c 1 10.0.3.3
  PING 10.0.3.3: 56  data bytes, press CTRL_C to break
    Reply from 10.0.3.3: bytes=56 Sequence=1 ttl=254 time=74 ms

  --- 10.0.3.3 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 74/74/74 ms

[R1]ping -c 1 10.0.24.4
  PING 10.0.24.4: 56  data bytes, press CTRL_C to break
    Reply from 10.0.24.4: bytes=56 Sequence=1 ttl=254 time=34 ms

  --- 10.0.24.4 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 34/34/34 ms
```

步骤三.   配置将外部路由引入到 OSPF

将R5的Loopback0接口网段10.0.5.0/24引入到OSPF区域。使用默认配置进行路由引入。

```
[R5]ospf 1
[R5-ospf-1]import-route direct
```

配置完成后，在R1上查看该路由，并测试网络连通性。

```
[R1]display ip routing-table protocol ospf
```

```
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Public routing table : OSPF
         Destinations : 7        Routes : 7

OSPF routing table status : <Active>
         Destinations : 7        Routes : 7

Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface

       10.0.2.0/24  OSPF    10   1562        D   10.0.12.2       Serial1/0/0
       10.0.3.0/24  OSPF    10   3124        D   10.0.12.2       Serial1/0/0
       10.0.5.0/24  O_ASE   150  1           D   10.0.12.2       Serial1/0/0
      10.0.23.0/24  OSPF    10   3124        D   10.0.12.2       Serial1/0/0
      10.0.24.0/24  OSPF    10   1563        D   10.0.12.2       Serial1/0/0
      10.0.35.0/24  OSPF    10   4686        D   10.0.12.2       Serial1/0/0
      10.0.35.3/32  O_ASE   150  1           D   10.0.12.2       Serial1/0/0

OSPF routing table status : <Inactive>
         Destinations : 0        Routes : 0

[R1]ping -c 1 10.0.5.5
  PING 10.0.5.5: 56  data bytes, press CTRL_C to break
    Reply from 10.0.5.5: bytes=56 Sequence=1 ttl=253 time=111 ms

  --- 10.0.5.5 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 111/111/111 ms
```

在R4上配置缺省路由，下一跳为Loopback0接口。并将该缺省路由引入到OSPF区域，定义为类型1，代价值为20，不使用永久发布。

```
[R4]ip route-static 0.0.0.0 0.0.0.0 LoopBack 0
[R4]ospf 1
[R4-ospf-1]default-route-advertise type 1 cost 20
[R4-ospf-1]quit
```

配置完成后，在R1上查看该缺省路由学习的情况。并测试网络的连通性。

```
[R1]display ip routing-table protocol ospf
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
```

```
Public routing table : OSPF
        Destinations : 7       Routes : 7

OSPF routing table status : <Active>
        Destinations : 7       Routes : 7

Destination/Mask    Proto   Pre  Cost     Flags NextHop        Interface

       0.0.0.0/0   O_ASE   150  1583        D   10.0.12.2      Serial1/0/0
      10.0.2.0/24  OSPF    10   1562        D   10.0.12.2      Serial1/0/0
      10.0.3.0/24  OSPF    10   3124        D   10.0.12.2      Serial1/0/0
      10.0.5.0/24  O_ASE   150  1           D   10.0.12.2      Serial1/0/0
     10.0.23.0/24  OSPF    10   3124        D   10.0.12.2      Serial1/0/0
     10.0.24.0/24  OSPF    10   1563        D   10.0.12.2      Serial1/0/0
     10.0.35.0/24  OSPF    10   4686        D   10.0.12.2      Serial1/0/0
     10.0.35.3/32  O_ASE   150  1           D   10.0.12.2      Serial1/0/0

OSPF routing table status : <Inactive>
        Destinations : 0       Routes : 0

[R1]ping -c 1 10.0.4.4
  PING 10.0.4.4: 56  data bytes, press CTRL_C to break
    Reply from 10.0.4.4: bytes=56 Sequence=1 ttl=254 time=39 ms

  --- 10.0.4.4 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 39/39/39 ms
```

步骤四.　配置区域 2 为 Stub 区域

在R1上查看路由信息。注意刚才看到的默认路由是外部路由（O_ASE），是通过R4发布的第五类LSA学习到的。

```
 [R1]display ospf lsdb

        OSPF Process 1 with Router ID 10.0.1.1
                Link State Database

                        Area: 0.0.0.2
 Type     LinkState ID    AdvRouter          Age  Len   Sequence   Metric
```

```
Router    10.0.2.2     10.0.2.2          12  48    80000003    1562
Router    10.0.1.1     10.0.1.1          11  60    80000003       0
Sum-Net   10.0.35.0    10.0.2.2          33  28    80000001    3124
Sum-Net   10.0.24.0    10.0.2.2          33  28    80000001       1
Sum-Net   10.0.3.0     10.0.2.2          33  28    80000001    1562
Sum-Net   10.0.2.0     10.0.2.2          33  28    80000001       0
Sum-Net   10.0.23.0    10.0.2.2          34  28    80000001    1562
Sum-Asbr  10.0.4.4     10.0.2.2          34  28    80000001       1
Sum-Asbr  10.0.5.5     10.0.2.2          34  28    80000001    3124


              AS External Database
Type      LinkState ID    AdvRouter          Age  Len   Sequence   Metric
External  0.0.0.0      10.0.4.4        1049  36    80000002      20
External  10.0.5.0     10.0.5.5        1350  36    80000001       1
External  10.0.35.0    10.0.5.5        1350  36    80000001       1
External  10.0.35.3    10.0.5.5        1350  36    80000001       1

[R1]display ospf lsdb ase 0.0.0.0

       OSPF Process 1 with Router ID 10.0.1.1
              Link State Database


 Type        : External
 Ls id       : 0.0.0.0
 Adv rtr     : 10.0.4.4
 Ls age      : 504
 Len         : 36
 Options     :  E
 seq#        : 80000002
 chksum      : 0xa981
 Net mask    : 0.0.0.0
 TOS 0  Metric: 20
 E type      : 1
 Forwarding Address : 0.0.0.0
 Tag         : 1
 Priority    : Low
```

在R1和R2上配置区域2为Stub区域。

```
[R1]ospf 1
[R1-ospf-1]area 2
[R1-ospf-1-area-0.0.0.2]stub
```

```
[R1-ospf-1-area-0.0.0.2]quit
[R1-ospf-1]quit

[R2]ospf 1
[R2-ospf-1]area 2
[R2-ospf-1-area-0.0.0.2]stub
[R2-ospf-1-area-0.0.0.2]quit
[R2-ospf-1]quit
```

配置完成后，在R1上对比之前的路由表，查看路由信息学习情况。这时可以看到，刚才的外部路由消失了，默认路由也变成了内部路由。

```
[R1]display ip routing-table protocol ospf
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Public routing table : OSPF
        Destinations : 6        Routes : 6

OSPF routing table status : <Active>
        Destinations : 6        Routes : 6

Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface

       0.0.0.0/0    OSPF    10   1563        D   10.0.12.2       Serial1/0/0
      10.0.2.0/24   OSPF    10   1562        D   10.0.12.2       Serial1/0/0
      10.0.3.0/24   OSPF    10   3124        D   10.0.12.2       Serial1/0/0
     10.0.23.0/24   OSPF    10   3124        D   10.0.12.2       Serial1/0/0
     10.0.24.0/24   OSPF    10   1563        D   10.0.12.2       Serial1/0/0
     10.0.35.0/24   OSPF    10   4686        D   10.0.12.2       Serial1/0/0

OSPF routing table status : <Inactive>
        Destinations : 0        Routes : 0
```

查看R1的LSDB，可以看到描述外部路由的LSA也消失了，默认路由是由一条第三类LSA学习到的。

```
[R1]display ospf lsdb

        OSPF Process 1 with Router ID 10.0.1.1
                Link State Database

                        Area: 0.0.0.2
 Type       LinkState ID    AdvRouter          Age  Len   Sequence   Metric
 Router     10.0.2.2        10.0.2.2           182  48    80000003    1562
```

```
 Router    10.0.1.1       10.0.1.1          182  60    80000004       0
 Sum-Net   0.0.0.0        10.0.2.2          183  28    80000001       1
 Sum-Net   10.0.35.0      10.0.2.2          183  28    80000001    3124
 Sum-Net   10.0.24.0      10.0.2.2          183  28    80000001       1
 Sum-Net   10.0.3.0       10.0.2.2          183  28    80000001    1562
 Sum-Net   10.0.2.0       10.0.2.2          184  28    80000001       0
 Sum-Net   10.0.23.0      10.0.2.2          184  28    80000001    1562
```

查看这条LSA的详细信息，可以发现这条默认路由是由R2发布的，这就验证了将一个区域配置为Stub区域以后，ABR会阻断第四、五类LSA向该区域发送，并通过三类LSA向该区域内泛洪一条默认路由指向ABR自己。

```
[R1]display ospf lsdb summary 0.0.0.0

       OSPF Process 1 with Router ID 10.0.1.1
                    Area: 0.0.0.2
             Link State Database


  Type       : Sum-Net
  Ls id      : 0.0.0.0
  Adv rtr    : 10.0.2.2
  Ls age     : 114
  Len        : 28
  Options    : None
  seq#       : 80000001
  chksum     : 0x1f31
  Net mask   : 0.0.0.0
  Tos 0  metric: 1
  Priority   : Low
```

在R2上将区域2配置为**no-summary**的完全Stub区域。

```
[R2]ospf 1
[R2-ospf-1]area 2
[R2-ospf-1-area-0.0.0.2]stub no-summary
[R2-ospf-1-area-0.0.0.2]quit
[R2-ospf-1]quit
```

查看R1的路由表，这时发现通过OSPF学习到的路由条目只剩一条默认路由了。

```
[R1]display ip routing-table protocol ospf
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
```

```
Public routing table : OSPF
        Destinations : 1        Routes : 1

OSPF routing table status : <Active>
        Destinations : 1        Routes : 1

Destination/Mask    Proto   Pre  Cost      Flags NextHop        Interface

       0.0.0.0/0    OSPF    10   1563        D   10.0.12.2      Serial1/0/0

OSPF routing table status : <Inactive>
        Destinations : 0        Routes : 0
```

查看R1的LSDB信息，发现除了R1和R2产生的第一类LSA以外，只剩一条由R2发布的三类LSA。

验证了在完全Stub区域中ABR会阻断了第三、四、五类LSA，并生成一条三类LSA，通告一条指向自己的默认路由。

```
[R1]display ospf lsdb

        OSPF Process 1 with Router ID 10.0.1.1
               Link State Database

                      Area: 0.0.0.2
 Type      LinkState ID    AdvRouter          Age  Len   Sequence   Metric
 Router    10.0.2.2        10.0.2.2           167  48    80000004    1562
 Router    10.0.1.1        10.0.1.1           166  60    80000006       0
 Sum-Net   0.0.0.0         10.0.2.2           549  28    80000001       1
```

步骤五. 配置区域 1 为 NSSA 区域

查看R3的路由表，R5发布的10.0.5.0/24是以外部路由的形式出现的。

```
[R3]display ip routing-table protocol ospf
Route Flags: R - relay, D - download to fib
----------------------------------------------------------------------------
Public routing table : OSPF
        Destinations : 7        Routes : 7

OSPF routing table status : <Active>
        Destinations : 6        Routes : 6
```

```
Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface

        0.0.0.0/0   O_ASE   150  1583        D   10.0.23.2       Serial2/0/0
       10.0.1.0/24  OSPF    10   3124        D   10.0.23.2       Serial2/0/0
       10.0.2.0/24  OSPF    10   1562        D   10.0.23.2       Serial2/0/0
       10.0.5.0/24  O_ASE   150  1           D   10.0.35.5       Serial3/0/0
      10.0.12.0/24  OSPF    10   3124        D   10.0.23.2       Serial2/0/0
      10.0.24.0/24  OSPF    10   1563        D   10.0.23.2       Serial2/0/0

OSPF routing table status : <Inactive>
         Destinations : 1        Routes : 1

Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface

      10.0.35.3/32  O_ASE   150  1               10.0.35.5       Serial3/0/0
```

在查看R5的路由表和LSDB信息。R5从R4学习到一条外部路由，其余的路由均是内部路由。R5通过第五类LSA向外发布了网络10.0.5.0/24。

```
[R5]display ip routing-table protocol ospf
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Public routing table : OSPF
         Destinations : 7        Routes : 7

OSPF routing table status : <Active>
         Destinations : 7        Routes : 7

Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface

        0.0.0.0/0   O_ASE   150  3145      D   10.0.35.3       Serial1/0/0
       10.0.1.0/24  OSPF    10   4686      D   10.0.35.3       Serial1/0/0
       10.0.2.0/24  OSPF    10   3124      D   10.0.35.3       Serial1/0/0
       10.0.3.0/24  OSPF    10   1562      D   10.0.35.3       Serial1/0/0
      10.0.12.0/24  OSPF    10   4686      D   10.0.35.3       Serial1/0/0
      10.0.23.0/24  OSPF    10   3124      D   10.0.35.3       Serial1/0/0
      10.0.24.0/24  OSPF    10   3125      D   10.0.35.3       Serial1/0/0

OSPF routing table status : <Inactive>
         Destinations : 0        Routes : 0

[R5]display ospf lsdb

         OSPF Process 1 with Router ID 10.0.5.5
```

```
              Link State Database

                     Area: 0.0.0.1
 Type      LinkState ID    AdvRouter          Age  Len   Sequence   Metric
 Router    10.0.5.5        10.0.5.5           882  48    80000004    1562
 Router    10.0.3.3        10.0.3.3          1309  48    80000003    1562
 Sum-Net   10.0.24.0       10.0.3.3            65  28    80000003    1563
 Sum-Net   10.0.12.0       10.0.3.3           819  28    80000001    3124
 Sum-Net   10.0.3.0        10.0.3.3            65  28    80000003       0
 Sum-Net   10.0.2.0        10.0.3.3            65  28    80000003    1562
 Sum-Net   10.0.1.0        10.0.3.3           812  28    80000001    3124
 Sum-Net   10.0.23.0       10.0.3.3            65  28    80000003    1562
 Sum-Asbr  10.0.4.4        10.0.3.3           602  28    80000002    1563


              AS External Database
 Type      LinkState ID    AdvRouter          Age  Len   Sequence   Metric
 External  10.0.5.0        10.0.5.5           882  36    80000002       1
 External  10.0.35.0       10.0.5.5           883  36    80000002       1
 External  10.0.35.3       10.0.5.5           883  36    80000002       1
 External  0.0.0.0         10.0.4.4           586  36    80000003      20
```

配置R3和R5的区域1为NSSA区域。

```
[R3]ospf
[R3-ospf-1]area 1
[R3-ospf-1-area-0.0.0.1]nssa

[R5]ospf
[R5-ospf-1]area 1
[R5-ospf-1-area-0.0.0.1]nssa
```

待邻居关系重新建立后，在R3上查看路由表。

```
[R3]display ip routing-table protocol ospf
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Public routing table : OSPF
        Destinations : 7        Routes : 7

OSPF routing table status : <Active>
        Destinations : 6        Routes : 6

Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface
```

```
      0.0.0.0/0   O_ASE    150  1583        D   10.0.23.2       Serial2/0/0
     10.0.1.0/24  OSPF     10   3124        D   10.0.23.2       Serial2/0/0
     10.0.2.0/24  OSPF     10   1562        D   10.0.23.2       Serial2/0/0
     10.0.5.0/24  O_NSSA   150  1           D   10.0.35.5       Serial3/0/0
    10.0.12.0/24  OSPF     10   3124        D   10.0.23.2       Serial2/0/0
    10.0.24.0/24  OSPF     10   1563        D   10.0.23.2       Serial2/0/0

OSPF routing table status : <Inactive>
        Destinations : 1        Routes : 1

Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface

    10.0.35.3/32  O_NSSA  150  1              10.0.35.5      Serial3/0/0
```

此时，R5通告的外部路由是以O_NSSA的形式出现在路由表里的。

再查看R5的路由表。

```
[R5]display ip routing-table protocol ospf
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Public routing table : OSPF
        Destinations : 7        Routes : 7

OSPF routing table status : <Active>
        Destinations : 7        Routes : 7

Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface

      0.0.0.0/0   O_NSSA  150  1           D   10.0.35.3       Serial1/0/0
     10.0.1.0/24  OSPF    10   4686        D   10.0.35.3       Serial1/0/0
     10.0.2.0/24  OSPF    10   3124        D   10.0.35.3       Serial1/0/0
     10.0.3.0/24  OSPF    10   1562        D   10.0.35.3       Serial1/0/0
    10.0.12.0/24  OSPF    10   4686        D   10.0.35.3       Serial1/0/0
    10.0.23.0/24  OSPF    10   3124        D   10.0.35.3       Serial1/0/0
    10.0.24.0/24  OSPF    10   3125        D   10.0.35.3       Serial1/0/0

OSPF routing table status : <Inactive>
        Destinations : 0        Routes : 0
```

刚才默认路由是以外部路由（O_ASE）的形式出现的，现在该默认路由变成了NSSA区域的外部路由（O_NSSA）。

查看R5的LSDB。

```
[R5]display ospf lsdb

         OSPF Process 1 with Router ID 10.0.5.5
                 Link State Database

                         Area: 0.0.0.1
 Type      LinkState ID     AdvRouter          Age  Len   Sequence   Metric
 Router    10.0.5.5         10.0.5.5           811  48    80000007    1562
 Router    10.0.3.3         10.0.3.3           811  48    80000007    1562
 Sum-Net   10.0.24.0        10.0.3.3           929  28    80000005    1563
 Sum-Net   10.0.12.0        10.0.3.3           929  28    80000005    3124
 Sum-Net   10.0.3.0         10.0.3.3           929  28    80000005       0
 Sum-Net   10.0.2.0         10.0.3.3           929  28    80000005    1562
 Sum-Net   10.0.1.0         10.0.3.3           930  28    80000005    3124
 Sum-Net   10.0.23.0        10.0.3.3           930  28    80000005    1562
 NSSA      10.0.5.0         10.0.5.5           819  36    80000005       1
 NSSA      10.0.35.0        10.0.5.5           819  36    80000006       1
 NSSA      10.0.35.3        10.0.5.5           819  36    80000005       1
 NSSA      0.0.0.0          10.0.3.3           930  36    80000005       1
```

发现，刚才的第五类LSA都消失了，外部路由以第七类LSA的形式向外通告。

查看默认路由的明细信息。

```
[R5]display ospf lsdb nssa 0.0.0.0

         OSPF Process 1 with Router ID 10.0.5.5
                         Area: 0.0.0.1
                 Link State Database


  Type      : NSSA
  Ls id     : 0.0.0.0
  Adv rtr   : 10.0.3.3
  Ls age    : 1149
  Len       : 36
  Options   : None
  seq#      : 80000005
  chksum    : 0x7745
  Net mask  : 0.0.0.0
  TOS 0  Metric: 1
  E type    : 2
  Forwarding Address : 0.0.0.0
  Tag       : 1
```

```
 Priority  : Low
```

刚才R5上的默认路由是R4通告给它的，而现在这条默认路由的通告者是R3。

从上面的结果我们可知，NSSA区域阻断了外部的第四、五类LSA进入，并且ABR会以第七类LSA的形式，向区域内通告一条默认路由。本区域的外部路由会以第七类LSA的形式，由ASBR向NSSA区域内通告。

NSSA和Stub区域的根本区别是，NSSA区域允许引入外部路由，而Stub区域不可以。

步骤六.　观察 NSSA 给 OSPF 带来的变化

使用**display ospf brief**命令查看R3目前所处的OSPF角色，可以看到，在Border Router这个字段有三个值：AREA　AS　NSSA。AREA表示该路由器是一台ABR，AS表示该路由器是一台ASBR，NSSA表示该路由器至少有一个接口位于NSSA区域。

```
[R3]display ospf brief

       OSPF Process 1 with Router ID 10.0.3.3
              OSPF Protocol Information

 RouterID: 10.0.3.3        Border Router:  AREA  AS  NSSA
 Multi-VPN-Instance is not enabled
 Global DS-TE Mode: Non-Standard IETF Mode
 Graceful-restart capability: disabled
 Helper support capability  : not configured
 Spf-schedule-interval: max 10000ms, start 500ms, hold 1000ms
 Default ASE parameters: Metric: 1 Tag: 1 Type: 2
 Route Preference: 10
 ASE Route Preference: 150
 SPF Computation Count: 15
 RFC 1583 Compatible
 Retransmission limitation is disabled
 Area Count: 2   Nssa Area Count: 1
 ExChange/Loading Neighbors: 0

Area: 0.0.0.0
 Authtype: None   Area flag: Normal
 SPF scheduled Count: 7
 ExChange/Loading Neighbors: 0

 Interface: 10.0.3.3 (LoopBack0)
```

```
Cost: 0      State: P-2-P    Type: P2P     MTU: 1500
Timers: Hello 10 , Dead 40 , Poll  120 , Retransmit 5 , Transmit Delay 1

 Interface: 10.0.23.3 (Serial2/0/0) --> 10.0.23.2
Cost: 1562   State: P-2-P    Type: P2P     MTU: 1500
Timers: Hello 10 , Dead 40 , Poll  120 , Retransmit 5 , Transmit Delay 1

Area: 0.0.0.1
Authtype: None   Area flag:   NSSA
SPF scheduled Count: 8
ExChange/Loading Neighbors: 0
NSSA Translator State: Elected

 Interface: 10.0.35.3 (Serial3/0/0) --> 10.0.35.5
Cost: 1562   State: P-2-P    Type: P2P     MTU: 1500
Timers: Hello 10 , Dead 40 , Poll  120 , Retransmit 5 , Transmit Delay 1
```

在NSSA区域中，由于不允许第五类LSA存在，所以ASBR是以第七类LSA的形式，向区域内通告外部路由的。但第七类LSA仅允许在NSSA区域内泛洪，NSSA区域的ABR收到这个第七类的LSA后，会将该第七类LSA转换成第五类LSA，然后向其他普通区域发布。

接下来我们在R3上观察7类LSA与5类LSA的转换过程。以10.0.5.0/24为例观察路由信息的传递。对于第七类LSA，Ls id描述了目的网段，Net mask描述了目的网段对应的掩码。Options字段为NP表示该LSA可以被ABR转化成一条第五类LSA，如果Options字段显示此LSA不可以被转换成第五类LSA，则Forwarding Address可以被设置成0.0.0.0；如果Options字段显示此LSA可以被转换成第五类LSA，则Forwarding Address不能被设置0.0.0.0。

在这里，所引入外部路由的下一跳不在OSPF路由域内，Forwarding Address设置为该ASBR上某个OSPF路由域内的Stub网段的接口IP地址。这里使用的地址为R5的Serial 1/0/0的接口地址。

```
[R3]display ospf lsdb nssa 10.0.5.0

        OSPF Process 1 with Router ID 10.0.3.3
                   Area: 0.0.0.0
             Link State Database

                   Area: 0.0.0.1
             Link State Database


  Type      : NSSA
```

```
  Ls id     : 10.0.5.0
  Adv rtr   : 10.0.5.5
  Ls age    : 836
  Len       : 36
  Options   :  NP
  seq#      : 80000001
  chksum    : 0xb0c2
  Net mask  : 255.255.255.0
  TOS 0  Metric: 1
  E type    : 2
  Forwarding Address : 10.0.35.5
  Tag       : 1
  Priority  : Low
```

查看R3生成的用于描述10.0.5.0/24的第五类LSA。

```
[R3]display ospf lsdb ase 10.0.5.0

       OSPF Process 1 with Router ID 10.0.3.3
              Link State Database


  Type      : External
  Ls id     : 10.0.5.0
  Adv rtr   : 10.0.3.3
  Ls age    : 882
  Len       : 36
  Options   :  E
  seq#      : 80000001
  chksum    : 0x413e
  Net mask  : 255.255.255.0
  TOS 0  Metric: 1
  E type    : 2
  Forwarding Address : 10.0.35.5
  Tag       : 1
  Priority  : Low
```

Ls id、Network Mask和Forwarding Address这几个字段的值直接从原来第七类LSA中拷贝。这样，10.0.5.0/24这个网段就被通告到其他区域了。

附加实验：思考并验证

NSSA区域类型适合用在哪些场景？

分析为什么R3路由器被定义为ASBR？

最终设备配置

```
<R1>display current-configuration
[V200R001C00SPC200]
#
 sysname R1
#
interface Serial1/0/0
 link-protocol ppp
 ip address 10.0.12.1 255.255.255.0
#
interface LoopBack0
 ip address 10.0.1.1 255.255.255.0
 ospf network-type broadcast
#
ospf 1 router-id 10.0.1.1
 area 0.0.0.2
  network 10.0.12.1 0.0.0.0
  network 10.0.1.1 0.0.0.0
  stub
#
return
```

```
<R2>display current-configuration
[V200R001C00SPC200]
#
 sysname R2
#
interface Serial1/0/0
 link-protocol ppp
 ip address 10.0.12.2 255.255.255.0
#
interface Serial2/0/0
 link-protocol ppp
 ip address 10.0.23.2 255.255.255.0
#
interface GigabitEthernet0/0/0
 ip address 10.0.24.2 255.255.255.0
#
interface LoopBack0
 ip address 10.0.2.2 255.255.255.0
```

```
 ospf network-type broadcast
#
ospf 1 router-id 10.0.2.2
 area 0.0.0.0
  network 10.0.23.2 0.0.0.0
  network 10.0.2.2 0.0.0.0
 area 0.0.0.2
  network 10.0.12.2 0.0.0.0
  stub no-summary
 area 0.0.0.3
  network 10.0.24.2 0.0.0.0
#
return

<R3>display current-configuration
[V200R001C00SPC200]
#
 sysname R3
#
interface Serial2/0/0
 link-protocol ppp
 ip address 10.0.23.3 255.255.255.0
#
interface Serial3/0/0
 link-protocol ppp
 ip address 10.0.35.3 255.255.255.0
#
interface LoopBack0
 ip address 10.0.3.3 255.255.255.0
 ospf network-type broadcast
#
ospf 1 router-id 10.0.3.3
 area 0.0.0.0
  network 10.0.23.3 0.0.0.0
  network 10.0.3.3 0.0.0.0
 area 0.0.0.1
  network 10.0.35.3 0.0.0.0
  nssa
#
return

<R4>display current-configuration
[V200R001C00SPC200]
```

```
#
 sysname R4
#
interface GigabitEthernet0/0/0
 ip address 10.0.24.4 255.255.255.0
#
interface NULL0
#
interface LoopBack0
 ip address 10.0.4.4 255.255.255.0
 ospf network-type broadcast
#
ospf 1 router-id 10.0.4.4
 default-route-advertise cost 20 type 1
 area 0.0.0.3
  network 10.0.24.4 0.0.0.0
#
 ip route-static 0.0.0.0 0.0.0.0 LoopBack0
#
return
```

```
<R5>display current-configuration
[V200R001C00SPC200]
#
 sysname R5
#
interface Serial1/0/0
 link-protocol ppp
 ip address 10.0.35.5 255.255.255.0
#
interface LoopBack0
 ip address 10.0.5.5 255.255.255.0
#
ospf 1 router-id 10.0.5.5
 import-route direct
 area 0.0.0.1
  network 10.0.35.5 0.0.0.0
  nssa
#
return
```

实验 2-5 OSPF 虚电路和区域路由过滤

学习目的

- 掌握使用OSPF虚电路来连接不连续的区域0的配置方法
- 掌握使用OSPF虚电路将非骨干区域连接到区域0的配置方法
- 掌握区域之间进行路由过滤和路由控制的方法

拓扑图

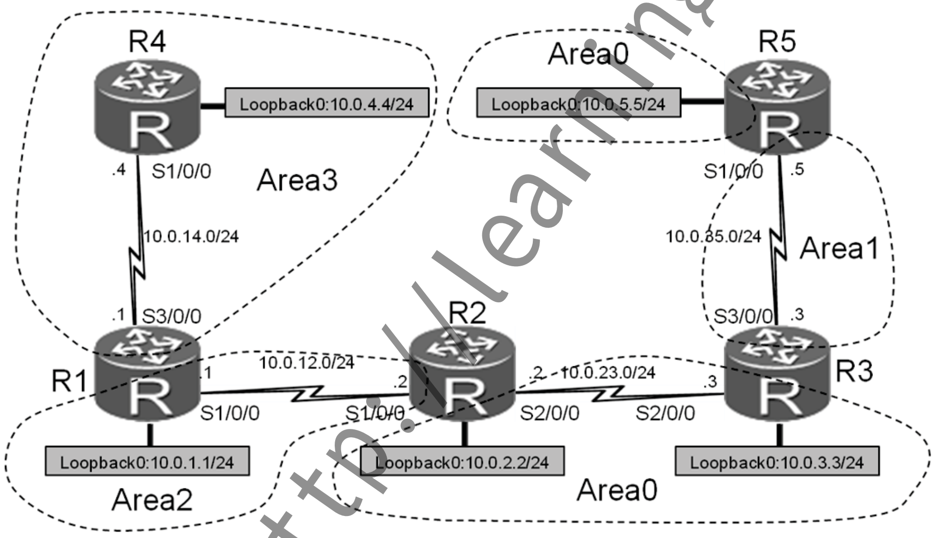

图2-5 OSPF虚电路和区域路由过滤

场景

你是公司的网络管理员。公司最近收购了两家小公司，他们的路由器是R4和R5。为了尽快合并网络，你决定先不去重新规划网络，而是使用虚电路实现网络互联。

网络直接相连后，你发现存在不连续的区域0，另外区域3与区域0没有直接连接。所以你决定在R1和R2之间建立虚电路，实现区域3与区域0的直接连接。另外在R3和R5之间建立虚电路，将不连续的区域0连接到一块。

同时为了明确设备的Router-ID，你配置设备使用固定的地址作为Router-ID。

学习任务

步骤一. 基础配置与 IP 编址

给所有路由器配置IP地址和掩码。配置时注意所有的Loopback接口配置掩码均为24位，模拟成一个单独的网段。

```
<R1>system-view
Enter system view, return user view with Ctrl+Z.
[R1]interface Serial 3/0/0
[R1-Serial3/0/0]ip address 10.0.14.1 24
[R1-Serial3/0/0]interface Serial 1/0/0
[R1-Serial1/0/0]ip address 10.0.12.1 24
[R1-Serial1/0/0]interface LoopBack 0
[R1-LoopBack0]ip address 10.0.1.1 24

<R2>system-view
Enter system view, return user view with Ctrl+Z.
[R2]interface Serial 1/0/0
[R2-Serial1/0/0]ip address 10.0.12.2 24
[R2-Serial1/0/0]interface Serial 2/0/0
[R2-Serial2/0/0]ip address 10.0.23.2 24
[R2-Serial2/0/0]interface LoopBack 0
[R2-LoopBack0]ip address 10.0.2.2 24

<R3>system-view
Enter system view, return user view with Ctrl+Z.
[R3]interface Serial 2/0/0
[R3-Serial2/0/0]ip address 10.0.23.3 24
[R3-Serial2/0/0]interface Serial 3/0/0
[R3-Serial3/0/0]ip address 10.0.35.3 24
[R3-Serial3/0/0]interface LoopBack 0
[R3-LoopBack0]ip address 10.0.3.3 24

<R4>system-view
Enter system view, return user view with Ctrl+Z.
[R4]interface Serial 1/0/0
[R4-Serial1/0/0]ip address 10.0.14.4 24
[R4-Serial1/0/0]interface LoopBack 0
```

```
[R4-LoopBack0]ip address 10.0.4.4 24

<R5>system-view
Enter system view, return user view with Ctrl+Z.
[R5]interface Serial 1/0/0
[R5-Serial1/0/0]ip address 10.0.35.5 24
[R5-Serial1/0/0]interface LoopBack 0
[R5-LoopBack0]ip address 10.0.5.5 24
```

配置完成后，测试直连链路的连通性。

```
[R1]ping -c 1 10.0.14.4
  PING 10.0.14.4: 56  data bytes, press CTRL_C to break
    Reply from 10.0.14.4: bytes=56 Sequence=1 ttl=255 time=39 ms

  --- 10.0.14.4 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 39/39/39 ms

[R1]ping -c 1 10.0.12.2
  PING 10.0.12.2: 56  data bytes, press CTRL_C to break
    Reply from 10.0.12.2: bytes=56 Sequence=1 ttl=255 time=29 ms

  --- 10.0.12.2 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 29/29/29 ms

[R3]ping -c 1 10.0.23.2
  PING 10.0.23.2: 56  data bytes, press CTRL_C to break
    Reply from 10.0.23.2: bytes=56 Sequence=1 ttl=255 time=45 ms

  --- 10.0.23.2 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 45/45/45 ms

[R3]ping -c 1 10.0.35.5
  PING 10.0.35.5: 56  data bytes, press CTRL_C to break
    Reply from 10.0.35.5: bytes=56 Sequence=1 ttl=255 time=32 ms
```

```
--- 10.0.35.5 ping statistics ---
  1 packet(s) transmitted
  1 packet(s) received
  0.00% packet loss
  round-trip min/avg/max = 32/32/32 ms
```

步骤二. 配置多区域 OSPF

在R1上配置Serial 1/0/0及Loopback 0属于区域2，Serial 3/0/0属于区域3。并对所有OSPF区域的Loopback接口，修改其OSPF网络类型为Broadcast类型，以便于OSPF发布Loopback口的真实掩码信息。所有的路由器使用Loopback 0的IP地址作为Router ID。

```
[R1]ospf 1 router-id 10.0.1.1
[R1-ospf-1]area 2
[R1-ospf-1-area-0.0.0.2]network 10.0.12.1 0.0.0.0
[R1-ospf-1-area-0.0.0.2]network 10.0.1.1 0.0.0.0
[R1-ospf-1-area-0.0.0.2]area 3
[R1-ospf-1-area-0.0.0.3]network 10.0.14.1 0.0.0.0
[R1-ospf-1-area-0.0.0.3]interface LoopBack 0
[R1-LoopBack0]ospf network-type broadcast
```

在R2上配置Serial 2/0/0及Loopback 0属于区域0，Serial 1/0/0属于区域2。

```
[R2]ospf 1 router-id 10.0.2.2
[R2-ospf-1]area 2
[R2-ospf-1-area-0.0.0.2]network 10.0.12.2 0.0.0.0
[R2-ospf-1-area-0.0.0.2]area 0
[R2-ospf-1-area-0.0.0.0]network 10.0.23.2 0.0.0.0
[R2-ospf-1-area-0.0.0.0]network 10.0.2.2 0.0.0.0
[R2-ospf-1-area-0.0.0.0]int LoopBack 0
[R2-LoopBack0]ospf network-type broadcast
```

在R3上配置Serial 2/0/0及Loopback 0属于区域0，Serial 3/0/0属于区域1。

```
[R3]ospf 1 router-id 10.0.3.3
[R3-ospf-1]area 0
[R3-ospf-1-area-0.0.0.0]network 10.0.23.3 0.0.0.0
[R3-ospf-1-area-0.0.0.0]network 10.0.3.3 0.0.0.0
[R3-ospf-1-area-0.0.0.0]area 1
[R3-ospf-1-area-0.0.0.1]network 10.0.35.3 0.0.0.0
```

```
[R3-ospf-1-area-0.0.0.1]interface LoopBack 0
[R3-LoopBack0]ospf network-type broadcast
```

在R4上配置Serial 1/0/0及Loopback 0属于区域3

```
[R4]ospf 1 router-id 10.0.4.4
[R4-ospf-1]area 3
[R4-ospf-1-area-0.0.0.3]network 10.0.14.4 0.0.0.0
[R4-ospf-1-area-0.0.0.3]network 10.0.4.4 0.0.0.0
[R4-ospf-1-area-0.0.0.3]interface LoopBack 0
[R4-LoopBack0]ospf network-type broadcast
```

在R5上配置Serial 1/0/0属于区域1，Loopback 0属于区域0。

```
[R5]ospf 1 router-id 10.0.5.5
[R5-ospf-1]area 0
[R5-ospf-1-area-0.0.0.0]network 10.0.5.5 0.0.0.0
[R5-ospf-1-area-0.0.0.0]area 1
[R5-ospf-1-area-0.0.0.1]network 10.0.35.5 0.0.0.0
[R5-ospf-1-area-0.0.0.1]interface LoopBack 0
[R5-LoopBack0]ospf network-type broadcast
```

步骤三. 查看每个路由器的路由表

查看R4的路由表，R4虽然与R1建立了邻接关系，没有学习到任何OSPF路由。

```
[R4]display ip routing-table protocol ospf
[R4]display ospf peer

       OSPF Process 1 with Router ID 10.0.4.4
             Neighbors

 Area 0.0.0.3 interface 10.0.14.4(Serial1/0/0)'s neighbors
 Router ID: 10.0.1.1        Address: 10.0.14.1
   State: Full  Mode:Nbr is  Slave  Priority: 1
   DR: None   BDR: None   MTU: 0
   Dead timer due in 39  sec
   Retrans timer interval: 4
   Neighbor is up for 00:21:33
   Authentication Sequence: [ 0 ]
```

再查看R4的LSDB发现仅存在第一类LSA，也就是说R1没有将其他区域的路由通告进区域3。

```
[R4]display ospf lsdb

        OSPF Process 1 with Router ID 10.0.4.4
               Link State Database

                      Area: 0.0.0.3
 Type      LinkState ID    AdvRouter          Age  Len   Sequence   Metric
 Router    10.0.4.4        10.0.4.4           571  60    80000005       0
 Router    10.0.1.1        10.0.1.1           616  48    80000003    1562
```

查看R1的路由表，缺失了10.0.5.0/24。至于缺少这条路由的原因，我们分析完R3的LSDB就明白了。

```
[R1]display ip routing-table protocol ospf
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Public routing table : OSPF
        Destinations : 5       Routes : 5

OSPF routing table status : <Active>
        Destinations : 5       Routes : 5

Destination/Mask   Proto   Pre  Cost       Flags NextHop         Interface

      10.0.2.0/24  OSPF    10   1562        D   10.0.12.2       Serial1/0/0
      10.0.3.0/24  OSPF    10   3124        D   10.0.12.2       Serial1/0/0
      10.0.4.0/24  OSPF    10   1562        D   10.0.14.4       Serial3/0/0
     10.0.23.0/24  OSPF    10   3124        D   10.0.12.2       Serial1/0/0
     10.0.35.0/24  OSPF    10   4686        D   10.0.12.2       Serial1/0/0

OSPF routing table status : <Inactive>
        Destinations : 0       Routes : 0
```

下面我们先来看一下R1的LSDB。为了避免区域间的环路，OSPF规定不允许直接在两个非骨干区域之间发布路由信息。从LSDB的角度来看，可以发现，ABR不会转发从非骨干区域收到的第三类LSA。

在R1上我们可以看到，在区域2的LSDB中有4条区域间路由，该路由是从R2（10.0.2.2）上学习到的，R1并没有将这些LSA转发到区域3里，所以R4学习不到非本区域的路由。

ABR也不会将从非骨干区域中学习到的路由转发给另一个非骨干区域，这里R1从R4这里学习到的路由不会以第三类LSA的形式通告进区域2，所以R2、R3、R5均学习不到区域3内的路由。

```
[R1]display ospf lsdb

         OSPF Process 1 with Router ID 10.0.1.1
                 Link State Database

                         Area: 0.0.0.2
 Type      LinkState ID    AdvRouter          Age  Len   Sequence   Metric
 Router    10.0.2.2        10.0.2.2          1251  48    80000023    1562
 Router    10.0.1.1        10.0.1.1          1266  60    80000024       0
 Sum-Net   10.0.35.0       10.0.2.2          1178  28    8000001B    3124
 Sum-Net   10.0.3.0        10.0.2.2          1178  28    8000001B    1562
 Sum-Net   10.0.2.0        10.0.2.2          1228  28    80000021       0
 Sum-Net   10.0.23.0       10.0.2.2          1189  28    8000001B    1562

                         Area: 0.0.0.3
 Type      LinkState ID    AdvRouter          Age  Len   Sequence   Metric
 Router    10.0.4.4        10.0.4.4           855  60    80000024       0
 Router    10.0.1.1        10.0.1.1           898  48    80000022    1562
```

查看R2的路由表，在R2的路由表中缺失了到达网络10.0.4.0/24、10.0.5.0/24、10.0.14.0/24的三条路由。

```
[R2]display ip routing-table protocol ospf
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Public routing table : OSPF
         Destinations : 3        Routes : 3

OSPF routing table status : <Active>
         Destinations : 3        Routes : 3

Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface

      10.0.1.0/24  OSPF    10   1562        D   10.0.12.1       Serial1/0/0
      10.0.3.0/24  OSPF    10   1562        D   10.0.23.3       Serial2/0/0
     10.0.35.0/24  OSPF    10   3124        D   10.0.23.3       Serial2/0/0

OSPF routing table status : <Inactive>
         Destinations : 0        Routes : 0
```

查看R2的LSDB，可以发现R1没有将区域3的路由通告给R2。

R2上会缺失到达网络10.0.4.0/24、10.0.14.0/24的路由。

在区域0中，R3也没有将10.0.5.0网络的路由通告给R2。

```
[R2]display ospf lsdb

	 OSPF Process 1 with Router ID 10.0.2.2
		 Link State Database

		         Area: 0.0.0.0
 Type      LinkState ID    AdvRouter          Age  Len   Sequence   Metric
 Router    10.0.3.3        10.0.3.3           973  60    80000027       0
 Router    10.0.2.2        10.0.2.2           972  60    80000028       0
 Sum-Net   10.0.35.0       10.0.3.3           984  28    8000001D    1562
 Sum-Net   10.0.12.0       10.0.2.2          1035  28    80000022    1562
 Sum-Net   10.0.1.0        10.0.2.2          1035  28    80000022    1562

		         Area: 0.0.0.2
 Type      LinkState ID    AdvRouter          Age  Len   Sequence   Metric
 Router    10.0.2.2        10.0.2.2          1046  48    80000024    1562
 Router    10.0.1.1        10.0.1.1          1063  60    80000025       0
 Sum-Net   10.0.35.0       10.0.2.2           973  28    8000001C    3124
 Sum-Net   10.0.3.0        10.0.2.2           973  28    8000001C    1562
 Sum-Net   10.0.2.0        10.0.2.2          1023  28    80000022       0
 Sum-Net   10.0.23.0       10.0.2.2           984  28    8000001C    1562
```

查看R3的路由表，缺失了到达网络10.0.4.0/24、10.0.5.0/24、10.0.14.0/24的路由。

```
[R3]display ip routing-table protocol ospf
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Public routing table : OSPF
         Destinations : 3        Routes : 3

OSPF routing table status : <Active>
         Destinations : 3        Routes : 3

Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface


      10.0.1.0/24  OSPF    10   3124        D   10.0.23.2       Serial2/0/0
      10.0.2.0/24  OSPF    10   1562        D   10.0.23.2       Serial2/0/0
     10.0.12.0/24  OSPF    10   3124        D   10.0.23.2       Serial2/0/0

OSPF routing table status : <Inactive>
         Destinations : 0        Routes : 0
```

查看R3的LSDB，我们可以发现，R3从区域1中收到了R5发来的第三类10.0.5.0。根据前面的原则，从非骨干区域收到的第三类LSA不会被转发。

R3没有将这条LSA再次发送到区域0中，这也正是R1和R2中没有10.0.5.0/24这条路由的原因。

```
[R3]display ospf lsdb

        OSPF Process 1 with Router ID 10.0.3.3
             Link State Database

                  Area: 0.0.0.0
 Type      LinkState ID    AdvRouter          Age  Len   Sequence   Metric
 Router    10.0.3.3        10.0.3.3           111  60    80000028       0
 Router    10.0.2.2        10.0.2.2           112  60    80000029       0
 Sum-Net   10.0.35.0       10.0.3.3           122  28    8000001E    1562
 Sum-Net   10.0.12.0       10.0.2.2           175  28    80000023    1562
 Sum-Net   10.0.1.0        10.0.2.2           175  28    80000023    1562

                  Area: 0.0.0.1
 Type      LinkState ID    AdvRouter          Age  Len   Sequence   Metric
 Router    10.0.5.5        10.0.5.5           117  48    8000001E    1562
 Router    10.0.3.3        10.0.3.3           117  48    80000020    1562
 Sum-Net   10.0.12.0       10.0.3.3           107  28    8000001D    3124
 Sum-Net   10.0.3.0        10.0.3.3           128  28    8000001D       0
 Sum-Net   10.0.2.0        10.0.3.3           107  28    8000001D    1562
 Sum-Net   10.0.1.0        10.0.3.3           108  28    8000001D    3124
 Sum-Net   10.0.5.0        10.0.5.5           128  28    8000001D       0
 Sum-Net   10.0.23.0       10.0.3.3           124  28    8000001D    1562
```

在这里我们注意到，R3的LSDB中已经有了R5发来的用于描述10.0.5.0/24的第三类LSA，但在R3的路由表上并没有出现这条路由。

查看R5的路由表。

```
[R5]display ip routing-table protocol ospf
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Public routing table : OSPF
        Destinations : 5        Routes : 5

OSPF routing table status : <Active>
        Destinations : 5        Routes : 5

Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface
```

```
    10.0.1.0/24  OSPF    10   4686        D   10.0.35.3       Serial1/0/0
    10.0.2.0/24  OSPF    10   3124        D   10.0.35.3       Serial1/0/0
    10.0.3.0/24  OSPF    10   1562        D   10.0.35.3       Serial1/0/0
   10.0.12.0/24  OSPF    10   4686        D   10.0.35.3       Serial1/0/0
   10.0.23.0/24  OSPF    10   3124        D   10.0.35.3       Serial1/0/0

OSPF routing table status : <Inactive>
       Destinations : 0      Routes : 0

 [R5]display ospf lsdb

       OSPF Process 1 with Router ID 10.0.5.5
              Link State Database

                    Area: 0.0.0.0
 Type      LinkState ID    AdvRouter         Age  Len   Sequence   Metric
 Router    10.0.5.5       10.0.5.5          820   36   80000002      0
 Sum-Net   10.0.35.0       10.0.5.5          861  28    80000001    1562

                    Area: 0.0.0.1
 Type      LinkState ID    AdvRouter         Age  Len   Sequence   Metric
 Router    10.0.5.5       10.0.5.5         1096   48   80000003    1562
 Router    10.0.3.3       10.0.3.3         1097   48   80000002    1562
 Sum-Net   10.0.12.0       10.0.3.3         1129  28    80000001    3124
 Sum-Net   10.0.3.0       10.0.3.3         1129  28    80000001      0
 Sum-Net   10.0.2.0       10.0.3.3         1129  28    80000001    1562
 Sum-Net   10.0.1.0       10.0.3.3         1129  28    80000001    3124
 Sum-Net   10.0.5.0       10.0.5.5          861  28    80000001      0
 Sum-Net   10.0.23.0       10.0.3.3         1129  28    80000001    1562
```

R5缺失了到达网络10.0.4.0/24、10.0.14.0/24的路由。

同时可以看到，R5上存在到达R3 Loopback 0的路由。

分析原因可知，R3存在连接到区域0的物理接口，即可以与区域0中其他的路由器交互路由信息 ,这时在R3上不会将非骨干区域发来的第三类LSA学习到的路由加到路由表里。R5虽然有接口在区域0中，但该接口为Loopback接口，Loopback接口在OSPF进行路由计算时为StubNet的链路类型。

查看R3的产生的类型一的LSA。

```
[R3]display ospf lsdb router 10.0.3.3

       OSPF Process 1 with Router ID 10.0.3.3
```

```
                 Area: 0.0.0.0
          Link State Database


 Type      : Router
 Ls id     : 10.0.3.3
 Adv rtr   : 10.0.3.3
 Ls age    : 732
 Len       : 60
 Options   :  ABR  E
 seq#      : 80000158
 chksum    : 0xde39
 Link count: 3
  * Link ID: 10.0.3.3
    Data    : 255.255.255.255
    Link Type: StubNet
    Metric : 0
    Priority : Medium
  * Link ID: 10.0.2.2
    Data    : 10.0.23.3
    Link Type: P-2-P
    Metric : 1562
  * Link ID: 10.0.23.0
    Data    : 255.255.255.0
    Link Type: StubNet
    Metric : 1562
    Priority : Low
```

从上面的输出中可以看到，R3与R2相连的链路的类型为P-2-P。类型为P-2-P、TransNet和Virtual类型的链路，路由器均认为该接口与其他路由器会交互路由信息。对于存在这三类链路连接到骨干区域的路由器不会将非骨干区域发来的第三类LSA加到路由表中。

```
[R5]display ospf lsdb router 10.0.5.5

        OSPF Process 1 with Router ID 10.0.5.5
                 Area: 0.0.0.0
          Link State Database


 Type      : Router
 Ls id     : 10.0.5.5
 Adv rtr   : 10.0.5.5
 Ls age    : 583
```

```
Len       : 36
Options    :  ABR  E
seq#      : 80000040
chksum     : 0x6d69
Link count: 1
 * Link ID: 10.0.5.5
   Data   : 255.255.255.255
   Link Type: StubNet
   Metric : 0
   Priority : Medium
```

在R5上，骨干区域中仅有一个Loopback 0，在描述这个接口的LSA中，链路的类型是StubNet，即末节网络，表示该接口不会再连接其他路由器，这时，路由器就会采用从非骨干区域发来的第三类LSA。

步骤四. 将不连续的区域 0 连在一块

在R3和R5上配置虚电路，注意在配置虚电路的时候在**vlink-peer**中配置的是对端ABR的Router ID。

```
[R3]ospf 1
[R3-ospf-1]area 1
[R3-ospf-1-area-0.0.0.1]vlink-peer 10.0.5.5

[R5]ospf
[R5-ospf-1]area 1
[R5-ospf-1-area-0.0.0.1]vlink-peer 10.0.3.3
```

然后检查虚电路邻居的状态是否为Full。

```
[R3]display ospf vlink

       OSPF Process 1 with Router ID 10.0.3.3
              Virtual Links

 Virtual-link Neighbor-id  -> 10.0.5.5, Neighbor-State: Full

 Interface: 10.0.35.3 (Serial3/0/0)
 Cost: 1562  State: P-2-P  Type: Virtual
 Transit Area: 0.0.0.1
 Timers: Hello 10 , Dead 40 , Retransmit 5 , Transmit Delay 1
 GR State: Normal
```

观察路由信息发生的变化。

```
[R3]display ip routing-table protocol ospf
Route Flags: R - relay, D - download to fib
----------------------------------------------------------------------------
Public routing table : OSPF
        Destinations : 4       Routes : 4

OSPF routing table status : <Active>
        Destinations : 4       Routes : 4

Destination/Mask    Proto  Pre  Cost      Flags NextHop         Interface

      10.0.1.0/24  OSPF    10   3124        D   10.0.23.2       Serial2/0/0
      10.0.2.0/24  OSPF    10   1562        D   10.0.23.2       Serial2/0/0
      10.0.5.0/24  OSPF    10   1562        D   10.0.35.5       Serial3/0/0
     10.0.12.0/24  OSPF    10   3124        D   10.0.23.2       Serial2/0/0

OSPF routing table status : <Inactive>
        Destinations : 0       Routes : 0
```

在R3上查看路由表发现已经学习到了10.0.5.0/24这条路由。

测试网络的连通性，R3可以与R5的Loopback 0连接的网段通讯。

```
[R3]ping -c 1 10.0.5.5
  PING 10.0.5.5: 56  data bytes, press CTRL_C to break
    Reply from 10.0.5.5: bytes=56 Sequence=1 ttl=255 time=34 ms

  --- 10.0.5.5 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 34/34/34 ms
```

查看R3的LSDB。

```
<R3>dis ospf lsdb

        OSPF Process 1 with Router ID 10.0.3.3
                Link State Database

                        Area: 0.0.0.0
 Type       LinkState ID    AdvRouter          Age  Len   Sequence   Metric
 Router      10.0.5.5        10.0.5.5          1098  48    80000005       0
```

```
Router    10.0.3.3        10.0.3.3          1096  72    80000008       0
Router    10.0.2.2        10.0.2.2           920  60    80000006       0
Sum-Net   10.0.35.0       10.0.3.3           830  28    80000002    1562
Sum-Net   10.0.35.0       10.0.5.5           565  28    80000002    1562
Sum-Net   10.0.12.0       10.0.2.2          1124  28    80000002    1562
Sum-Net   10.0.1.0        10.0.2.2          1110  28    80000002    1562

                    Area: 0.0.0.1
Type      LinkState ID    AdvRouter           Age  Len   Sequence   Metric
Router    10.0.5.5        10.0.5.5          1098  48    80000004    1562
Router    10.0.3.3        10.0.3.3          1096  48    80000003    1562
Sum-Net   10.0.12.0       10.0.3.3           830  28    80000002    3124
Sum-Net   10.0.3.0        10.0.3.3           831  28    80000002       0
Sum-Net   10.0.2.0        10.0.3.3           831  28    80000002    1562
Sum-Net   10.0.1.0        10.0.3.3           831  28    80000002    3124
Sum-Net   10.0.5.0        10.0.5.5           566  28    80000002       0
Sum-Net   10.0.23.0       10.0.3.3           831  28    80000002    1562
```

看到在R3上共收到了2条来自R5的第一类LSA。第一条是在区域0中收到的，虚电路属于区域0，所以该LSA实际上是通过虚电路学习到的。另一条第一类LSA是在区域1中学习到的，这条LSA刚才没建虚电路的时候就有。路由表中的10.0.5.0/24路由是通过区域0学习到的LSA计算出的。

查看R3的LSDB中关于10.0.5.5这条第一类LSA的详细信息。

```
[R3]display ospf lsdb router 10.0.5.5

       OSPF Process 1 with Router ID 10.0.3.3
                    Area: 0.0.0.0
              Link State Database


 Type      : Router
 Ls id     : 10.0.5.5
 Adv rtr    : 10.0.5.5
 Ls age    : 621
 Len        : 48
 Options    :  ABR  E
 seq#       : 80000005
 chksum     : 0x1291
 Link count: 2
  * Link ID: 10.0.5.0
    Data    : 255.255.255.0
    Link Type: StubNet
    Metric : 0
```

```
    Priority : Low
   * Link ID: 10.0.3.3
    Data   : 10.0.35.5
    Link Type: Virtual
    Metric : 1562
                 Area: 0.0.0.1
             Link State Database


  Type      : Router
  Ls id     : 10.0.5.5
  Adv rtr    : 10.0.5.5
  Ls age     : 621
  Len       : 48
  Options    :  ABR  VIRTUAL  E
  seq#      : 80000004
  chksum     : 0x3530
  Link count: 2
   * Link ID: 10.0.3.3
    Data   : 10.0.35.5
    Link Type: P-2-P
    Metric : 1562
   * Link ID: 10.0.35.0
    Data   : 255.255.255.0
    Link Type: StubNet
    Metric : 1562
    Priority : Low
```

可以看到这条LSA中描述了网络10.0.5.0/24，所以在R3上就有了这条路由。而从区域1中学习到的这条第一类LSA仅描述了R3和R5的互联网段。

查看R5的LSDB。

```
[R5]display ospf lsdb

        OSPF Process 1 with Router ID 10.0.5.5
             Link State Database

                 Area: 0.0.0.0
 Type      LinkState ID    AdvRouter          Age  Len   Sequence   Metric
 Router    10.0.5.5       10.0.5.5          577   48    80000005      0
 Router    10.0.3.3       10.0.3.3          577   72    80000008      0
 Router    10.0.2.2       10.0.2.2          401   60    80000006      0
 Sum-Net   10.0.35.0      10.0.5.5           45   28    80000002    1562
```

```
 Sum-Net   10.0.35.0       10.0.3.3          312  28    80000002    1562
 Sum-Net   10.0.12.0       10.0.2.2          606  28    80000002    1562
 Sum-Net   10.0.1.0        10.0.2.2          593  28    80000002    1562

                     Area: 0.0.0.1
 Type      LinkState ID    AdvRouter          Age  Len   Sequence   Metric
 Router    10.0.5.5        10.0.5.5          578  48    80000004    1562
 Router    10.0.3.3        10.0.3.3          578  48    80000003    1562
 Sum-Net   10.0.12.0       10.0.3.3          313  28    80000002    3124
 Sum-Net   10.0.3.0        10.0.3.3          313  28    80000002       0
 Sum-Net   10.0.2.0        10.0.3.3          313  28    80000002    1562
 Sum-Net   10.0.1.0        10.0.3.3          313  28    80000002    3124
 Sum-Net   10.0.5.0        10.0.5.5           46  28    80000002       0
 Sum-Net   10.0.23.0       10.0.3.3          313  28    80000002    1562
```

可以发现和R3的LSDB是一样的。建立的虚电路以后，R3和R5均有接口属于区域0了，所以LSDB是同步的。

步骤五. 区域 3 通过虚电路连接到区域 0

在R1和R2上配置虚电路。

```
[R1]ospf 1
[R1-ospf-1]area 2
[R1-ospf-1-area-0.0.0.2]vlink-peer 10.0.2.2
[R1-ospf-1-area-0.0.0.2]quit
[R1-ospf-1]quit

[R2]ospf
[R2-ospf-1]area 2
[R2-ospf-1-area-0.0.0.2]vlink-peer 10.0.1.1
[R2-ospf-1-area-0.0.0.2]quit
[R2-ospf-1]quit
```

查看R4的OSPF路由表。

```
[R4]display ip routing-table protocol ospf
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Public routing table : OSPF
         Destinations : 7        Routes : 7
```

```
OSPF routing table status : <Active>
        Destinations : 7       Routes : 7

Destination/Mask    Proto   Pre  Cost      Flags NextHop        Interface

      10.0.1.0/24  OSPF    10   1562       D   10.0.14.1       Serial1/0/0
      10.0.2.0/24  OSPF    10   3124       D   10.0.14.1       Serial1/0/0
      10.0.3.0/24  OSPF    10   4686       D   10.0.14.1       Serial1/0/0
      10.0.5.0/24  OSPF    10   6248       D   10.0.14.1       Serial1/0/0
     10.0.12.0/24  OSPF    10   3124       D   10.0.14.1       Serial1/0/0
     10.0.23.0/24  OSPF    10   4686       D   10.0.14.1       Serial1/0/0
     10.0.35.0/24  OSPF    10   6248       D   10.0.14.1       Serial1/0/0

OSPF routing table status : <Inactive>
        Destinations : 0       Routes : 0
```

发现该路由器已拥有了全网路由。

测试网络的连通性。

```
[R4]ping -c 1 10.0.5.5
  PING 10.0.5.5: 56  data bytes, press CTRL_C to break
    Reply from 10.0.5.5: bytes=56 Sequence=1 ttl=252 time=132 ms

  --- 10.0.5.5 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 132/132/132 ms
```

查看R1的LSDB。

```
[R1]display ospf lsdb

        OSPF Process 1 with Router ID 10.0.1.1
                Link State Database

                    Area: 0.0.0.0
 Type      LinkState ID    AdvRouter          Age  Len   Sequence   Metric
 Router    10.0.5.5        10.0.5.5           419  48    80000006       0
 Router    10.0.3.3        10.0.3.3           418  72    80000009       0
 Router    10.0.2.2        10.0.2.2           232  72    8000000A       0
 Router    10.0.1.1        10.0.1.1           233  36    80000001    1562
 Sum-Net   10.0.35.0       10.0.3.3           151  28    80000003    1562
 Sum-Net   10.0.35.0       10.0.5.5          1687  28    80000002    1562
```

```
Sum-Net   10.0.14.0      10.0.1.1          291  28    80000001     1562
Sum-Net   10.0.12.0      10.0.1.1          291  28    80000001     1562
Sum-Net   10.0.12.0      10.0.2.2          444  28    80000003     1562
Sum-Net   10.0.1.0       10.0.1.1          291  28    80000001        0
Sum-Net   10.0.1.0       10.0.2.2          430  28    80000003     1562
Sum-Net   10.0.4.0       10.0.1.1          291  28    80000001     1562

                 Area: 0.0.0.2
Type      LinkState ID    AdvRouter          Age  Len   Sequence   Metric
Router    10.0.2.2       10.0.2.2          235  48    80000005     1562
Router    10.0.1.1       10.0.1.1          234  60    80000009        0
Sum-Net   10.0.35.0      10.0.2.2          151  28    80000003     3124
Sum-Net   10.0.14.0      10.0.1.1          291  28    80000001     1562
Sum-Net   10.0.3.0       10.0.2.2          234  28    80000003     1562
Sum-Net   10.0.2.0       10.0.2.2          443  28    80000003        0
Sum-Net   10.0.5.0       10.0.2.2          402  28    80000002     3124
Sum-Net   10.0.4.0       10.0.1.1          292  28    80000001     1562
Sum-Net   10.0.23.0      10.0.2.2          286  28    80000003     1562

                 Area: 0.0.0.3
Type      LinkState ID    AdvRouter          Age  Len   Sequence   Metric
Router    10.0.4.4       10.0.4.4         1193  60    80000005        0
Router    10.0.1.1       10.0.1.1          292  48    80000004     1562
Sum-Net   10.0.35.0      10.0.1.1          292  28    80000001     4686
Sum-Net   10.0.12.0      10.0.1.1          294  28    80000001     1562
Sum-Net   10.0.3.0       10.0.1.1          294  28    80000001     3124
Sum-Net   10.0.2.0       10.0.1.1          294  28    80000001     1562
Sum-Net   10.0.1.0       10.0.1.1          294  28    80000001        0
Sum-Net   10.0.5.0       10.0.1.1          294  28    80000001     4686
Sum-Net   10.0.23.0      10.0.1.1          294  28    80000001     3124
```

由于创建了虚电路，R1中有了区域0的LSA，这样区域0和区域3就可以直接交互路由信息了。R1把区域0中的路由信息以第三类LSA的形式通告进了区域3。

查看R4的LSDB。

```
[R4]display ospf lsdb

        OSPF Process 1 with Router ID 10.0.4.4
                 Link State Database

                 Area: 0.0.0.3
Type      LinkState ID    AdvRouter          Age  Len   Sequence   Metric
Router    10.0.4.4       10.0.4.4         1303  60    80000005        0
```

```
Router    10.0.1.1      10.0.1.1         404  48    80000004    1562
Sum-Net   10.0.35.0     10.0.1.1         404  28    80000001    4686
Sum-Net   10.0.12.0     10.0.1.1         404  28    80000001    1562
Sum-Net   10.0.3.0      10.0.1.1         404  28    80000001    3124
Sum-Net   10.0.2.0      10.0.1.1         404  28    80000001    1562
Sum-Net   10.0.1.0      10.0.1.1         405  28    80000001       0
Sum-Net   10.0.5.0      10.0.1.1         405  28    80000001    4686
Sum-Net   10.0.23.0     10.0.1.1         405  28    80000001    3124
```

可以看到该路由器学习到了R1发布的第三类LSA。

R4有其他区域的路由。

步骤六. 配置区域之间的路由过滤

控制10.0.4.0/24网段的路由信息的发布。使R1可以学到该路由，但R2、R3、R5学不到这条路由。

设置一个访问控制列表。

```
[R1]acl number 2000
[R1-acl-basic-2000]rule deny source 10.0.4.0 0.0.0.255
[R1-acl-basic-2000]rule permit
```

R1配置针对类型3的LSA的过滤，配置在区域3向其他区域发送更新时进行过滤。

```
[R1]ospf 1
[R1-ospf-1]area 3
[R1-ospf-1-area-0.0.0.3]filter 2000 export
```

在R2上查看路由信息过滤的情况。

```
[R2]dis ip routing-table protocol ospf
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Public routing table : OSPF
         Destinations : 5        Routes : 5

OSPF routing table status : <Active>
         Destinations : 5        Routes : 5

Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface

       10.0.1.0/24  OSPF    10   1562        D   10.0.12.1       Serial1/0/0
```

```
      10.0.3.0/24  OSPF   10   1562       D  10.0.23.3      Serial2/0/0
      10.0.5.0/24  OSPF   10   3124       D  10.0.23.3      Serial2/0/0
     10.0.14.0/24  OSPF   10   3124       D  10.0.12.1      Serial1/0/0
     10.0.35.0/24  OSPF   10   3124       D  10.0.23.3      Serial2/0/0

OSPF routing table status : <Inactive>
        Destinations : 0       Routes : 0
```

R2已经学习不到路由10.0.4.0/24了。

此时R1的路由表中仍然有该条目。因为R1和R4是同一个区域的，R4通过第一类LSA把该路由通告给R1。

```
[R1]display ip routing-table protocol ospf
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Public routing table : OSPF
        Destinations : 6       Routes : 6

OSPF routing table status : <Active>
        Destinations : 6       Routes : 6

Destination/Mask    Proto  Pre  Cost      Flags NextHop         Interface

      10.0.2.0/24  OSPF   10   1562       D  10.0.12.2      Serial1/0/0
      10.0.3.0/24  OSPF   10   3124       D  10.0.12.2      Serial1/0/0
      10.0.4.0/24  OSPF   10   1562       D  10.0.14.4      Serial3/0/0
      10.0.5.0/24  OSPF   10   4686       D  10.0.12.2      Serial1/0/0
     10.0.23.0/24  OSPF   10   3124       D  10.0.12.2      Serial1/0/0
     10.0.35.0/24  OSPF   10   4686       D  10.0.12.2      Serial1/0/0

OSPF routing table status : <Inactive>
        Destinations : 0       Routes : 0
```

附加实验：思考并验证

OSPF协议中为什么区域0必须连续？从当前的OSPF设计的角度来看，是否可以对类型1和类型2的LSA进行过滤？

最终设备配置

```
<R1>display current-configuration
```

```
[V200R001C00SPC200]
#
 sysname R1
#
acl number 2000
 rule 5 deny source 10.0.4.0 0.0.0.255
 rule 10 permit
#
interface Serial1/0/0
 link-protocol ppp
 ip address 10.0.12.1 255.255.255.0
#
interface Serial3/0/0
 link-protocol ppp
 ip address 10.0.14.1 255.255.255.0
#
interface LoopBack0
 ip address 10.0.1.1 255.255.255.0
 ospf network-type broadcast
#
ospf 1 router-id 10.0.1.1
 area 0.0.0.0
 area 0.0.0.2
  network 10.0.1.1 0.0.0.0
  network 10.0.12.1 0.0.0.0
  vlink-peer 10.0.2.2
 area 0.0.0.3
  filter 2000 export
  network 10.0.14.1 0.0.0.0
#
return

<R2>display current-configuration
[V200R001C00SPC200]
#
 sysname R2
#
interface Serial1/0/0
 link-protocol ppp
 ip address 10.0.12.2 255.255.255.0
#
interface Serial2/0/0
 link-protocol ppp
```

```
 ip address 10.0.23.2 255.255.255.0
#
interface LoopBack0
 ip address 10.0.2.2 255.255.255.0
 ospf network-type broadcast
#
ospf 1 router-id 10.0.2.2
 area 0.0.0.0
  network 10.0.23.2 0.0.0.0
  network 10.0.2.2 0.0.0.0
 area 0.0.0.2
  network 10.0.12.2 0.0.0.0
  vlink-peer 10.0.1.1
#
return

<R3>display current-configuration
[V200R001C00SPC200]
#
 sysname R3
#
interface Serial2/0/0
 link-protocol ppp
 ip address 10.0.23.3 255.255.255.0
#
interface Serial3/0/0
 link-protocol ppp
 ip address 10.0.35.3 255.255.255.0
#
interface LoopBack0
 ip address 10.0.3.3 255.255.255.0
 ospf network-type broadcast
#
ospf 1 router-id 10.0.3.3
 area 0.0.0.0
  network 10.0.3.3 0.0.0.0
  network 10.0.23.3 0.0.0.0
 area 0.0.0.1
  network 10.0.35.3 0.0.0.0
  vlink-peer 10.0.5.5
#
return
```

```
<R4>display current-configuration
[V200R001C00SPC200]
#
 sysname R4
#
interface Serial1/0/0
 link-protocol ppp
 ip address 10.0.14.4 255.255.255.0
#
interface LoopBack0
 ip address 10.0.4.4 255.255.255.0
 ospf network-type broadcast
#
ospf 1 router-id 10.0.4.4
 area 0.0.0.3
  network 10.0.14.4 0.0.0.0
  network 10.0.4.4 0.0.0.0
#
return
```

```
<R5>display current-configuration
[V200R001C00SPC200]
#
 sysname R5
#
interface Serial1/0/0
 link-protocol ppp
 ip address 10.0.35.5 255.255.255.0
#
interface LoopBack0
 ip address 10.0.5.5 255.255.255.0
 ospf network-type broadcast
#
ospf 1 router-id 10.0.5.5
 area 0.0.0.0
  network 10.0.5.5 0.0.0.0
 area 0.0.0.1
  network 10.0.35.5 0.0.0.0
  vlink-peer 10.0.3.3
#
return
```

实验 2-6 OSPF 故障排除

学习目的

- 掌握对单区域OSPF中区域号码不匹配进行故障排除的方法
- 掌握对单区域OSPF中掩码不匹配进行故障排除的方法
- 掌握对单区域OSPF中Hello时间不匹配进行故障排除的方法
- 掌握对单区域OSPF中Router-id冲突进行故障排除的方法
- 掌握OSPF认证相关的故障排除方法
- 掌握OSPF汇总相关的故障排除方法
- 掌握虚电路相关的故障排除方法

拓扑图

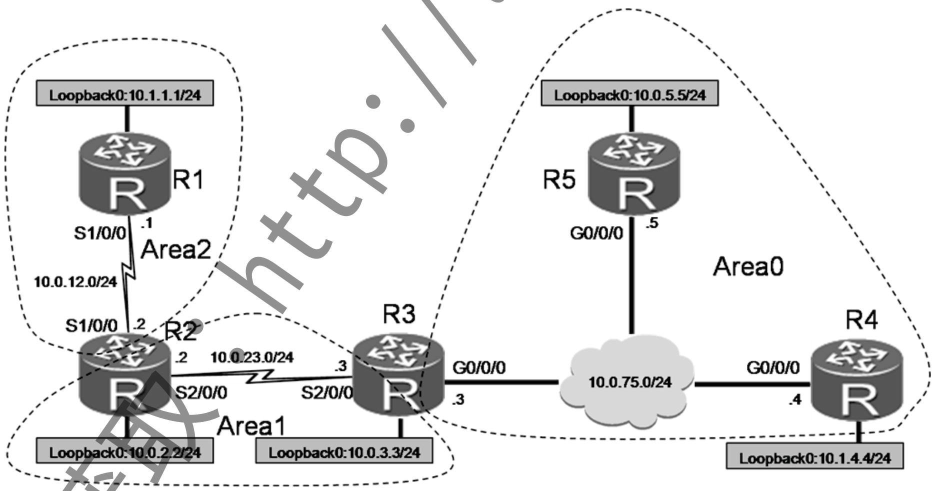

图2-6 OSPF 故障排除

场景

你是公司的网络管理员。公司的网络采用了OSPF协议作为路由协议。OSPF协议功能强大，但是相应的配置也较为复杂。并且在网络规划中，你使用了OSPF的各种特性，同时也使用了虚链路。在实施过程中，你碰到很多的网络通讯问题。不过庆幸的是，通过使用故障排除的思想和方法，你成功的找到了各种错误，并实现了网络的恢复。

学习任务

步骤一. 基础配置与 IP 编址

给所有路由器配置IP地址和掩码。配置时注意所有的Loopback接口配置掩码均为24位，模拟成一个单独的网段。

```
<R1>system-view
Enter system view, return user view with Ctrl+Z.
[R1]int Serial 1/0/0
[R1-Serial1/0/0]ip address 10.0.12.1 24
[R1-Serial1/0/0]interface LoopBack 0
[R1-LoopBack0]ip address 10.1.1.1 24

<R2>system-view
Enter system view, return user view with Ctrl+Z.
[R2]int Serial 1/0/0
[R2-Serial1/0/0]ip address 10.0.12.2 24
[R2-Serial1/0/0]interface Serial 2/0/0
[R2-Serial2/0/0]ip address 10.0.23.2 24
[R2-Serial2/0/0]interface LoopBack 0
[R2-LoopBack0]ip address 10.0.2.2 24
```

为模拟相应的错误，R3的G0/0/0接口配置IP地址为10.0.75.3/25，其余接口地址按照拓扑图中的标识进行配置。

```
<R3>system-view
Enter system view, return user view with Ctrl+Z.
[R3]interface Serial 2/0/0
[R3-Serial2/0/0]ip address 10.0.23.3 24
[R3-Serial2/0/0]interface GigabitEthernet 0/0/0
[R3-GigabitEthernet0/0/0]ip address 10.0.75.3 25
[R3-GigabitEthernet0/0/0]interface LoopBack 0
[R3-LoopBack0]ip address 10.0.3.3 24
```

```
<R4>system-view
Enter system view, return user view with Ctrl+Z.
[R4]interface GigabitEthernet 0/0/0
[R4-GigabitEthernet0/0/0]ip address 10.0.75.4 24
[R4-GigabitEthernet0/0/0]interface LoopBack 0
[R4-LoopBack0]ip address 10.1.4.4 24

<R5>system-view
Enter system view, return user view with Ctrl+Z.
[R5]interface GigabitEthernet 0/0/0
[R5-GigabitEthernet0/0/0]ip address 10.0.75.5 24
[R5-GigabitEthernet0/0/0]interface LoopBack 0
[R5-LoopBack0]ip address 10.0.5.5 24
```

配置完成后，测试直连链路的连通性。

```
[R3]ping -c 1 10.0.75.4
  PING 10.0.75.4: 56  data bytes, press CTRL_C to break
    Reply from 10.0.75.4: bytes=56 Sequence=1 ttl=255 time=5 ms

  --- 10.0.75.4 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 5/5/5 ms

[R3]ping -c 1 10.0.75.5
  PING 10.0.75.5: 56  data bytes, press CTRL_C to break
    Reply from 10.0.75.5: bytes=56 Sequence=1 ttl=255 time=5 ms

  --- 10.0.75.5 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 5/5/5 ms

[R3]ping -c 1 10.0.23.2
  PING 10.0.23.2: 56  data bytes, press CTRL_C to break
    Reply from 10.0.23.2: bytes=56 Sequence=1 ttl=255 time=41 ms

  --- 10.0.23.2 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
```

```
    0.00% packet loss
    round-trip min/avg/max = 41/41/41 ms

[R1]ping -c 1 10.0.12.2
  PING 10.0.12.2: 56  data bytes, press CTRL_C to break
    Reply from 10.0.12.2: bytes=56 Sequence=1 ttl=255 time=37 ms

  --- 10.0.12.2 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 37/37/37 ms
```

步骤二.　配置多区域 OSPF

配置R1的接口Serial 1/0/0和Loopback 0属于区域2 ,使用接口Loopback 0的地址作为Router ID。注意对所有OSPF区域的Loopback接口，修改其OSPF网络类型为Broadcast类型，以便于OSPF发布Loopback口的真实掩码信息。

```
[R1]ospf 1 router-id 10.1.1.1
[R1-ospf-1]area 2
[R1-ospf-1-area-0.0.0.2]network 10.0.12.1 0.0.0.0
[R1-ospf-1-area-0.0.0.2]network 10.1.1.1 0.0.0.0
[R1-ospf-1-area-0.0.0.2]interface LoopBack 0
[R1-LoopBack0]ospf network-type broadcast
```

在R2上配置接口Serial 2/0/0及Loopback 0属于区域1，接口Serial 1/0/0属于区域2，在启用OSPF时没有静态指定Router ID。

```
[R2]ospf 1
[R2-ospf-1]area 1
[R2-ospf-1-area-0.0.0.1]network 10.0.23.2 0.0.0.0
[R2-ospf-1-area-0.0.0.1]network 10.0.2.2 0.0.0.0
[R2-ospf-1-area-0.0.0.1]area 2
[R2-ospf-1-area-0.0.0.2]network 10.0.12.2 0.0.0.0
[R2-ospf-1-area-0.0.0.2]int LoopBack 0
[R2-LoopBack0]ospf network-type broadcast
```

在R3上配置接口Serial 2/0/0及Loopback 0属于区域1，接口GigabitEthernet 0/0/0属于区域0。

```
[R3]ospf 1 router-id 10.0.3.3
[R3-ospf-1]area 1
```

```
[R3-ospf-1-area-0.0.0.1]network 10.0.23.3 0.0.0.0
[R3-ospf-1-area-0.0.0.1]network 10.0.3.3 0.0.0.0
[R3-ospf-1-area-0.0.0.1]area 0
[R3-ospf-1-area-0.0.0.0]network 10.0.75.3 0.0.0.0
[R3-ospf-1-area-0.0.0.0]interface LoopBack 0
[R3-LoopBack0]ospf network-type broadcast
```

在R4上配置接口GigabitEthernet 0/0/0属于区域1，接口Loopback 0不属于任何区域。在配置OSPF进程时使用**ospf 1 router-id**指定R4的Router ID为10.0.5.5。

```
[R4]ospf 1 router-id 10.0.5.5
[R4-ospf-1]area 1
[R4-ospf-1-area-0.0.0.1]network 10.0.75.4 0.0.0.0
```

在R5上配置接口GigabitEthernet 0/0/0和Loopback 0属于区域0。

```
[R5]ospf 1 router-id 10.0.5.5
[R5-ospf-1]area 0
[R5-ospf-1-area-0.0.0.0]network 10.0.75.5 0.0.0.0
[R5-ospf-1-area-0.0.0.0]network 10.0.5.5 0.0.0.0
[R5-ospf-1-area-0.0.0.0]interface LoopBack 0
[R5-LoopBack0]ospf network-type broadcast
```

步骤三. 在区域内排除 OSPF 故障

查看R4邻居列表，发现R4没有与其他路由器建立邻居关系。

```
[R4]display ospf peer

       OSPF Process 1 with Router ID 10.0.5.5
```

在R3、R4、R5上分别运行**display ospf error**查看OSPF发生的错误。

```
[R3]display ospf error

       OSPF Process 1 with Router ID 10.0.3.3
              OSPF error statistics

General packet errors:
 0     : IP: received my own packet     2450  : Bad packet
 0     : Bad version                    0     : Bad checksum
 1032  : Bad area id                    0     : Drop on unnumbered interface
```

```
 0     : Bad virtual link              0     : Bad authentication type
 0     : Bad authentication key        0     : Packet too small
 0     : Packet size > ip length       0     : Transmit error
 2     : Interface down                0     : Unknown neighbor

HELLO packet errors:
 2450  : Netmask mismatch              0     : Hello timer mismatch
 0     : Dead timer mismatch           0     : Extern option mismatch
 0     : Router id confusion           0     : Virtual neighbor unknown
 0     : NBMA neighbor unknown         0     : Invalid Source Address

[R4]display ospf error

       OSPF Process 1 with Router ID 10.0.5.5
              OSPF error statistics

General packet errors:
 0     : IP: received my own packet 1354  : Bad packet
 0     : Bad version                   0     : Bad checksum
 0     : Bad area id                   0     : Drop on unnumbered interface
 1032  : Bad virtual link              0     : Bad authentication type
 0     : Bad authentication key        0     : Packet too small
 0     : Packet size > ip length       0     : Transmit error
 3     : Interface down                0     : Unknown neighbor

HELLO packet errors:
 1354  : Netmask mismatch              0     : Hello timer mismatch
 0     : Dead timer mismatch           0     : Extern option mismatch
 1155  : Router id confusion           0     : Virtual neighbor unknown
 0     : NBMA neighbor unknown         0     : Invalid Source Address

[R5]display ospf error

       OSPF Process 1 with Router ID 10.0.5.5
              OSPF error statistics

General packet errors:
 0     : IP: received my own packet    1216  : Bad packet
 0     : Bad version                   0     : Bad checksum
 0     : Bad area id                   0     : Drop on unnumbered interface
 0     : Bad virtual link              0     : Bad authentication type
 0     : Bad authentication key        0     : Packet too small
 0     : Packet size > ip length       0     : Transmit error
```

```
 3     : Interface down                   0     : Unknown neighbor

HELLO packet errors:
 1216  : Netmask mismatch                 0     : Hello timer mismatch
 0     : Dead timer mismatch              0     : Extern option mismatch
 1169  : Router id confusion              0     : Virtual neighbor unknown
 0     : NBMA neighbor unknown            0     : Invalid Source Address
```

从上面的输出中，我们可以看到，在R3、R4、R5这三台路由器之间一共发生了五种错误：Router ID冲突（Router id confusion）、子网掩码不匹配（Netmask mismatch）、错误的区域号（Bad area id）、错误的数据包（Bad packet）、错误的虚电路（Bad virtual link）。

在这里我们并没有配置虚电路，在这种情况下错误的虚电路其实就是区域号错误。对于R4来说，它在一个区域号为1的接口上收到了一个区域号为0的OSPF数据包，它会认为这是一个通过虚电路发送过来的数据包。路由器本身没有配置虚电路，就发生了这种错误。

子网掩码错误也是错误的数据包的一部分，我们可以先修正子网掩码的问题再观察是否还有错误的数据包。

首先排除Router ID冲突的问题。我们可以依次查看每台路由器的Router ID来手工找出发生冲突的路由器，也可以通过系统日志来查找。通过**display logbuffer**翻阅路由器当前的系统日志。

```
[R5]display logbuffer
Logging buffer configuration and contents: enabled
Allowed max buffer size: 1024
Actual buffer size: 512
Channel number: 4, Channel name: logbuffer
Dropped messages: 0
Overwritten messages: 0
Current messages: 66

Nov 29 2011 13:38:22+00:00 R5 %%01OSPF/4/CONFLICT_ROUTERID_INTF(l)[0]:OSPF
Router id conflict is detected on interface. (ProcessId=1, RouterId=10.0.5.5,
AreaId=0.0.0.0, InterfaceName=GigabitEthernet0/0/0,  IpAddr=10.0.75.5,
PacketSrcIp=10.0.75.4)
```

从上面可以看到，与R5发生冲突的路由器接口的IP地址是10.0.75.4。查看拓扑，发现10.0.75.4是R4的接口地址。然后查看R4的Router ID，可以看到该路由器的Router ID和R5是一样的。同时还发现R4的区域号配置也有误。

```
[R4]display ospf brief
```

```
        OSPF Process 1 with Router ID 10.0.5.5
              OSPF Protocol Information

 RouterID: 10.0.5.5       Border Router:
 Multi-VPN-Instance is not enabled
 Global DS-TE Mode: Non-Standard IETF Mode
 Graceful-restart capability: disabled
 Helper support capability  : not configured
 Spf-schedule-interval: max 10000ms, start 500ms, hold 1000ms
 Default ASE parameters: Metric: 1 Tag: 1 Type: 2
 Route Preference: 10
 ASE Route Preference: 150
 SPF Computation Count: 13
 RFC 1583 Compatible
 Retransmission limitation is disabled
 Area Count: 1   Nssa Area Count: 0
 ExChange/Loading Neighbors: 0

 Area: 0.0.0.1
 Authtype: None   Area flag: Normal
 SPF scheduled Count: 2
 ExChange/Loading Neighbors: 0

 Interface: 10.0.75.4 (GigabitEthernet0/0/0)
 Cost: 1      State: DR       Type: Broadcast    MTU: 1500
 Priority: 1
 Designated Router: 10.0.75.4
 Backup Designated Router: 0.0.0.0
 Timers: Hello 10 , Dead 40 , Poll  120 , Retransmit 5 , Transmit Delay 1
```

修改R4的Router ID和区域号。

```
[R4]ospf 1 router-id 10.1.4.4
[R4-ospf-1]area 0
[R4-ospf-1-area-0.0.0.0]network 10.0.75.4 0.0.0.0

<R4>reset ospf process
Warning: The OSPF process will be reset. Continue? [Y/N]:y
```

修改完成以后通过命令**reset ospf counter**清空OSPF计数器。

注意reset命令需在用户视图下运行。

```
<R4>reset ospf counters
```

重置后，稍等片刻，再运行display ospf error检查该问题是否消失。

```
<R4>display ospf error

        OSPF Process 1 with Router ID 10.1.4.4
               OSPF error statistics

General packet errors:
 0     : IP: received my own packet     2     : Bad packet
 0     : Bad version                    0     : Bad checksum
 0     : Bad area id                    0     : Drop on unnumbered interface
 0     : Bad virtual link               0     : Bad authentication type
 0     : Bad authentication key         0     : Packet too small
 0     : Packet size > ip length        0     : Transmit error
 0     : Interface down                 0     : Unknown neighbor

HELLO packet errors:
 2     : Netmask mismatch               0     : Hello timer mismatch
 0     : Dead timer mismatch            0     : Extern option mismatch
 0     : Router id confusion            0     : Virtual neighbor unknown
 0     : NBMA neighbor unknown          0     : Invalid Source Address
```

可以看到在修改完配置以后，Route ID冲突和区域号错误的问题消失了，还剩下子网掩码不匹配的问题。为了找出是哪台路由器配置了错误的子网掩码，我们在R4上查看Debug信息。

```
<R4>terminal debugging
Info: Current terminal debugging is on.
<R4>debugging ospf packet hello
Nov 29 2011 14:56:16.720.1+00:00 R4 RM/6/RMDEBUG:
 FileID: 0xd0178024 Line: 2178 Level: 0x20
 OSPF 1: RECV Packet. Interface: GigabitEthernet0/0/0
Nov 29 2011 14:56:16.720.2+00:00 R4 RM/6/RMDEBUG:  Source Address: 10.0.75.3
Nov 29 2011 14:56:16.720.3+00:00 R4 RM/6/RMDEBUG:  Destination Address: 224.0.0.5
Nov 29 2011 14:56:16.720.4+00:00 R4 RM/6/RMDEBUG:  Ver# 2, Type: 1 (Hello)
Nov 29 2011 14:56:16.720.5+00:00 R4 RM/6/RMDEBUG:  Length: 44, Router: 10.0.3.3
Nov 29 2011 14:56:16.720.6+00:00 R4 RM/6/RMDEBUG:  Area: 0.0.0.0, Chksum: 9a18
Nov 29 2011 14:56:16.720.7+00:00 R4 RM/6/RMDEBUG:  AuType: 00
Nov 29 2011 14:56:16.720.8+00:00 R4 RM/6/RMDEBUG:  Key(ascii): 0 0 0 0 0 0 0 0
Nov 29 2011 14:56:16.720.9+00:00 R4 RM/6/RMDEBUG:  Net Mask: 255.255.255.128
Nov 29 2011 14:56:16.720.10+00:00 R4 RM/6/RMDEBUG:  Hello Int: 10, Option: _E_
Nov 29 2011 14:56:16.720.11+00:00 R4 RM/6/RMDEBUG:  Rtr Priority: 1, Dead Int:
40
Nov 29 2011 14:56:16.720.12+00:00 R4 RM/6/RMDEBUG:  DR: 10.0.75.3
```

```
Nov 29 2011 14:56:16.720.13+00:00 R4 RM/6/RMDEBUG:  BDR: 0.0.0.0
Nov 29 2011 14:56:16.730.1+00:00 R4 RM/6/RMDEBUG:  # Attached Neighbors: 0
```

从上面的信息我们可以看出，从10.0.75.3发来的Hello包中子网掩码是255.255.255.128。查看拓扑，发现R3的对应接口配置错误。

```
[R3-GigabitEthernet0/0/0]display this
[V200R001C00SPC200]
#
interface GigabitEthernet0/0/0
 ip address 10.0.75.3 255.255.255.128
#
return
[R3-GigabitEthernet0/0/0]ip address 10.0.75.3 24
```

再次清空OSPF计数器，查看是否还存在错误。

```
<R3>reset ospf counters

<R3>display ospf error

        OSPF Process 1 with Router ID 10.0.3.3
               OSPF error statistics

General packet errors:
 0     : IP: received my own packet     0     : Bad packet
 0     : Bad version                    0     : Bad checksum
 0     : Bad area id                    0     : Drop on unnumbered interface
 0     : Bad virtual link               0     : Bad authentication type
 0     : Bad authentication key         0     : Packet too small
 0     : Packet size > ip length        0     : Transmit error
 0     : Interface down                 0     : Unknown neighbor

HELLO packet errors:
 0     : Netmask mismatch               0     : Hello timer mismatch
 0     : Dead timer mismatch            0     : Extern option mismatch
 0     : Router id confusion            0     : Virtual neighbor unknown
 0     : NBMA neighbor unknown          0     : Invalid Source Address
```

在R3上检查邻居列表，发现各邻居的状态已正常。

```
[R3]display ospf peer brief

        OSPF Process 1 with Router ID 10.0.3.3
                Peer Statistic Information
```

```
 ----------------------------------------------------------------------------
 Area Id          Interface                       Neighbor id      State
 0.0.0.0          GigabitEthernet0/0/0              10.1.4.4        Full
 0.0.0.0          GigabitEthernet0/0/0              10.0.5.5        Full
 0.0.0.1          Serial2/0/0                       10.0.2.2        Full
 ----------------------------------------------------------------------------
```

接下来我们修改R4的GigabitEthernet 0/0/0接口的Hello间隔为5秒，观察邻居关系是否可以形成。

```
[R4]interface GigabitEthernet 0/0/0
[R4-GigabitEthernet0/0/0]ospf timer hello 5
```

经过约半分钟以后，可以观察到R4的邻居都消失了。

```
[R4]display ospf peer brief

        OSPF Process 1 with Router ID 10.1.4.4
                Peer Statistic Information
 ----------------------------------------------------------------------------
 Area Id          Interface                       Neighbor id      State
 ----------------------------------------------------------------------------
```

查看OSPF的错误。

```
[R4]display ospf error

        OSPF Process 1 with Router ID 10.1.4.4
               OSPF error statistics

General packet errors:
 0     : IP: received my own packet     2     : Bad packet
 0     : Bad version                    0     : Bad checksum
 0     : Bad area id                    0     : Drop on unnumbered interface
 0     : Bad virtual link               0     : Bad authentication type
 0     : Bad authentication key         0     : Packet too small
 0     : Packet size > ip length        0     : Transmit error
 0     : Interface down                 0     : Unknown neighbor

HELLO packet errors:
 0     : Netmask mismatch               2     : Hello timer mismatch
 0     : Dead timer mismatch            0     : Extern option mismatch
 0     : Router id confusion            0     : Virtual neighbor unknown
 0     : NBMA neighbor unknown          0     : Invalid Source Address
```

可以看到有Hello时间不匹配的错误出现，说明OSPF要求邻居间Hello间隔一样。

取消Hello间隔的修改。再次检查邻居列表。

```
[R4-GigabitEthernet0/0/0]undo ospf timer hello
[R4]display ospf peer brief

       OSPF Process 1 with Router ID 10.1.4.4
              Peer Statistic Information
 ----------------------------------------------------------------------------
 Area Id         Interface                   Neighbor id      State
 0.0.0.0         GigabitEthernet0/0/0          10.0.3.3         Full
 0.0.0.0         GigabitEthernet0/0/0          10.0.5.5         Full
 ----------------------------------------------------------------------------
```

发现邻居关系已恢复正常。

步骤四.　OSPF 认证故障排除

在R1和R2上配置基于接口的认证。

其中R1采用simple方式，密钥为123。

R2采用MD5方式，密钥为huawei。

```
[R1]interface Serial 1/0/0
[R1-Serial1/0/0]ospf authentication-mode simple plain 123

[R2]interface Serial 1/0/0
[R2-Serial1/0/0]ospf authentication-mode md5 1 plain huawei
```

配置完成以后在R1上可以查看到OSPF的错误。

```
[R1-Serial1/0/0]display ospf error

       OSPF Process 1 with Router ID 10.1.1.1
             OSPF error statistics

General packet errors:
 0     : IP: received my own packet      15    : Bad packet
 0     : Bad version                     0     : Bad checksum
 0     : Bad area id                     0     : Drop on unnumbered interface
 0     : Bad virtual link                15    : Bad authentication type
 0     : Bad authentication key          0     : Packet too small
```

```
 0     : Packet size > ip length        0     : Transmit error
 0     : Interface down                 0     : Unknown neighbor
```

将R1的认证方式配置为MD5后，查看是否还存在错误。

```
[R1]interface Serial 1/0/0
[R1-Serial1/0/0]ospf authentication-mode md5 1 plain 123
[R1-Serial1/0/0]return
<R1>reset ospf counters
<R1>display ospf error

       OSPF Process 1 with Router ID 10.1.1.1
             OSPF error statistics

General packet errors:
 0     : IP: received my own packet     1     : Bad packet
 0     : Bad version                    0     : Bad checksum
 0     : Bad area id                    0     : Drop on unnumbered interface
 0     : Bad virtual link               1     : Bad authentication type
 0     : Bad authentication key         0     : Packet too small
 0     : Packet size > ip length        0     : Transmit error
 0     : Interface down                 0     : Unknown neighbor
```

可以看到该问题还存在。

将R1的密钥也改成huawei，观察邻居关系。

```
[R1]interface Serial 1/0/0
[R1-Serial1/0/0] ospf authentication-mode md5 1 plain huawei
[R1-Serial1/0/0]quit
[R1]display ospf peer brief

       OSPF Process 1 with Router ID 10.1.1.1
              Peer Statistic Information
 ----------------------------------------------------------------------------
 Area Id         Interface                   Neighbor id      State
 0.0.0.2         Serial1/0/0                  10.0.2.2        Full
 ----------------------------------------------------------------------------
```

可见，R1与R2已建立邻接关系。

步骤五. 虚电路故障排除

为保证区域2与区域0之间的连通性，在R2和R3之间创建虚电路。

```
[R2]ospf 1
[R2-ospf-1]area 1
[R2-ospf-1-area-0.0.0.1]vlink-peer 10.0.3.3

[R3]ospf 1
[R3-ospf-1]area 1
[R3-ospf-1-area-0.0.0.1]vlink-peer 10.0.2.2
```

观察虚电路建立是否正常，以及R1是否学习到了全网路由。

```
[R2]display ospf vlink

        OSPF Process 1 with Router ID 10.0.2.2
               Virtual Links

 Virtual-link Neighbor-id  -> 10.0.3.3, Neighbor-State: Full

 Interface: 10.0.23.2 (Serial2/0/0)
 Cost: 1562  State: P-2-P  Type: Virtual
 Transit Area: 0.0.0.1
 Timers: Hello 10 , Dead 40 , Retransmit 5 , Transmit Delay 1
 GR State: Normal

[R1]display ip routing-table protocol ospf
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Public routing table : OSPF
        Destinations : 5        Routes : 5

OSPF routing table status : <Active>
        Destinations : 5        Routes : 5

Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface

      10.0.2.0/24  OSPF    10   1562        D   10.0.12.2       Serial1/0/0
      10.0.3.0/24  OSPF    10   3124        D   10.0.12.2       Serial1/0/0
      10.0.5.0/24  OSPF    10   3125        D   10.0.12.2       Serial1/0/0
     10.0.23.0/24  OSPF    10   3124        D   10.0.12.2       Serial1/0/0
     10.0.75.0/24  OSPF    10   3125        D   10.0.12.2       Serial1/0/0

OSPF routing table status : <Inactive>
        Destinations : 0        Routes : 0
```

在R1上测试连通性，证实可以到达R5。

```
[R1]ping -c 1 10.0.5.5
  PING 10.0.5.5: 56  data bytes, press CTRL_C to break
    Reply from 10.0.5.5: bytes=56 Sequence=1 ttl=253 time=81 ms

  --- 10.0.5.5 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 81/81/81 ms
```

由于测试的需要，删除R2的loopback0接口。

```
[R2]undo interface LoopBack 0
```

后来由于一次偶然的事故，路由器重启了。在这里我们通过重启OSPF进程的方法模拟路由器重启。

```
<R2>reset ospf process
Warning: The OSPF process will be reset. Continue? [Y/N]:y
```

这时连接到R1的用户发现自己无法访问区域外的地址。管理员登录到R1上发现无法与R5的Loopback地址通讯。

```
[R1]ping -c 1 10.0.5.5
  PING 10.0.5.5: 56  data bytes, press CTRL_C to break
    Request time out

  --- 10.0.5.5 ping statistics ---
    1 packet(s) transmitted
    0 packet(s) received
    100.00% packet loss
```

检查R2和R3之间的虚电路之后发现状态不正常，同时发现R2的Router ID发生了变化。

```
[R2]display ospf vlink

         OSPF Process 1 with Router ID 10.0.23.2
                 Virtual Links

 Virtual-link Neighbor-id  -> 10.0.3.3, Neighbor-State: Down

 Interface: 10.0.23.2 (Serial2/0/0)
 Cost: 1562  State: P-2-P  Type: Virtual
 Transit Area: 0.0.0.1
```

```
 Timers: Hello 10 , Dead 40 , Retransmit 5 , Transmit Delay 1
 GR State: Normal
```

由于虚电路的建立是基于对端设备的Router ID的。R2的Router ID发生了变化，所以虚电路发生了故障。

通常我们在启动OSPF进程时指定该进程的Router ID，就是为了防止路由器在运行过程中Router ID发生变化。

下面我们将R2的Router ID固定为10.0.2.2，并将Loopback地址添加回去，然后重启OSPF进程。

```
[R2]ospf 1 router-id 10.0.2.2
Info: The configuration succeeded. You need to restart the OSPF process to validate
the new router ID.
[R2-ospf-1]interface LoopBack 0
[R2-LoopBack0]ip address 10.0.2.2 24
[R2-LoopBack0]quit
<R2>reset ospf process
Warning: The OSPF process will be reset. Continue? [Y/N]:y
```

再次查看虚电路状态。

```
[R2]display ospf vlink

        OSPF Process 1 with Router ID 10.0.2.2
               Virtual Links

 Virtual-link Neighbor-id  -> 10.0.3.3, Neighbor-State: Full

 Interface: 10.0.23.2 (Serial2/0/0)
 Cost: 1562  State: P-2-P  Type: Virtual
 Transit Area: 0.0.0.1
 Timers: Hello 10 , Dead 40 , Retransmit 5 , Transmit Delay 1
 GR State: Normal
```

此时虚电路已恢复正常。

管理员出于安全的考虑，在区域0使用了基于区域的认证，启用了MD5对报文进行加密，密钥为huawei。

```
[R3]ospf 1
[R3-ospf-1]area 0
[R3-ospf-1-area-0.0.0.0]authentication-mode md5 1 plain huawei

[R4]ospf 1
```

```
[R4-ospf-1]area 0
[R4-ospf-1-area-0.0.0.0]authentication-mode md5 1 plain huawei

[R5]ospf 1
[R5-ospf-1]area 0
[R5-ospf-1-area-0.0.0.0]authentication-mode md5 1 plain huawei
```

这时，管理员再次发现区域2中的用户无法访问区域外的网络，检查虚电路后发现虚电路又出于故障的状态。

```
[R2]dis ospf vlink

       OSPF Process 1 with Router ID 10.0.2.2
              Virtual Links

 Virtual-link Neighbor-id  -> 10.0.3.3, Neighbor-State: Down

 Interface: 10.0.23.2 (Serial2/0/0)
 Cost: 1562  State: P-2-P  Type: Virtual
 Transit Area: 0.0.0.1
 Timers: Hello 10 , Dead 40 , Retransmit 5 , Transmit Delay 1
```

检查OSPF的错误发现有认证错误发生。

```
[R2]display ospf error

       OSPF Process 1 with Router ID 10.0.2.2
              OSPF error statistics

General packet errors:
 0     : IP: received my own packet     2     : Bad packet
 0     : Bad version                    0     : Bad checksum
 0     : Bad area id                    0     : Drop on unnumbered interface
 0     : Bad virtual link               2     : Bad authentication type
 0     : Bad authentication key         0     : Packet too small
 0     : Packet size > ip length        0     : Transmit error
 0     : Interface down                 0     : Unknown neighbor

HELLO packet errors:
 0     : Netmask mismatch               0     : Hello timer mismatch
 0     : Dead timer mismatch            0     : Extern option mismatch
 0     : Router id confusion            0     : Virtual neighbor unknown
 0     : NBMA neighbor unknown          0     : Invalid Source Address
```

OSPF的虚电路属于区域0。区域0打开了基于区域的认证，虚电路上也需要打开认证。

```
[R2]ospf
[R2-ospf-1]area 0
[R2-ospf-1-area-0.0.0.0]authentication-mode md5 1 plain huawei
```

这时虚电路的状态恢复了正常，R1也能正常访问其他区域了。

```
[R2]display ospf vlink

        OSPF Process 1 with Router ID 10.0.2.2
               Virtual Links

 Virtual-link Neighbor-id  -> 10.0.3.3, Neighbor-State: Full

 Interface: 10.0.23.2 (Serial2/0/0)
 Cost: 1562  State: P-2-P  Type: Virtual
 Transit Area: 0.0.0.1
 Timers: Hello 10 , Dead 40 , Retransmit 5 , Transmit Delay 1
 GR State: Normal

[R1]ping -c 1 10.0.5.5
  PING 10.0.5.5: 56  data bytes, press CTRL_C to break
    Reply from 10.0.5.5: bytes=56 Sequence=1 ttl=253 time=73 ms

  --- 10.0.5.5 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 73/73/73 ms
```

步骤六. OSPF 路由汇总故障排除

首先在R4上以外部路由的形式引入Loopback 0接口地址，并进行地址汇总，汇总后的子网掩码为16位。

```
[R4]ospf
[R4-ospf-1]import-route direct
[R4-ospf-1]asbr-summary 10.1.0.0 255.255.0.0
```

一段时间之后，管理员在R2上配置了区域间汇总，将R1的Loopback 0接口连接的网段汇总成16位掩码的路由。

```
[R2]ospf
[R2-ospf-1]area 2
[R2-ospf-1-area-0.0.0.2]abr-summary 10.1.0.0 255.255.0.0
```

这时，除了连接到R4的用户以外，全网所有用户均反馈不能访问R4的Loopback地址10.1.4.4。

检查与R4同一区域的路由器R5的路由表发现，若要到达10.1.4.4，匹配到路由条目10.1.0.0/16，而该路由的下一跳是10.0.75.3。

为何会产生这样一个错误的条目呢？

```
[R5]display ip routing-table protocol ospf
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Public routing table : OSPF
         Destinations : 5        Routes : 5

OSPF routing table status : <Active>
         Destinations : 5        Routes : 5

Destination/Mask    Proto   Pre  Cost    Flags NextHop        Interface

      10.0.2.0/24  OSPF    10   1563      D   10.0.75.3    GigabitEthernet0/0/0
      10.0.3.0/24  OSPF    10   1         D   10.0.75.3    GigabitEthernet0/0/0
     10.0.12.0/24  OSPF    10   3125      D   10.0.75.3    GigabitEthernet0/0/0
     10.0.23.0/24  OSPF    10   1563      D   10.0.75.3    GigabitEthernet0/0/0
      10.1.0.0/16  OSPF    10   3125      D   10.0.75.3    GigabitEthernet0/0/0

OSPF routing table status : <Inactive>
         Destinations : 0        Routes : 0
```

我们先来观察一下R5的LSDB。

```
[R5]display ospf lsdb

         OSPF Process 1 with Router ID 10.0.5.5
                 Link State Database

                          Area: 0.0.0.0
 Type        LinkState ID    AdvRouter          Age  Len   Sequence   Metric
 Router        10.0.5.5       10.0.5.5           214  48    80000025       0
 Router        10.0.3.3       10.0.3.3          1246  48    80000024       1
 Router        10.0.2.2       10.0.2.2          1247  36    80000005    1562
 Router        10.1.4.4       10.1.4.4           648  36    8000000D       1
```

```
 Network        10.0.75.4      10.1.4.4         206  36   80000004      0
 Sum-Net    10.0.12.0      10.0.2.2         916  28   80000002    1562
 Sum-Net    10.0.3.0       10.0.3.3         893  28   80000008       0
 Sum-Net    10.0.3.0       10.0.2.2         916  28   80000002    1562
 Sum-Net    10.0.2.0       10.0.3.3         919  28   80000003    1562
 Sum-Net    10.0.2.0       10.0.2.2         916  28   80000002       0
 Sum-Net    10.1.0.0       10.0.2.2         538  28   80000001    1562
 Sum-Net    10.0.23.0      10.0.3.3         893  28   80000008    1562
 Sum-Net    10.0.23.0      10.0.2.2         917  28   80000002    1562


              AS External Database
 Type      LinkState ID   AdvRouter          Age  Len   Sequence   Metric
 External        10.0.75.0      10.1.4.4         649  36   80000001      1
 External        10.1.0.0       10.1.4.4         620  36   80000001      2
```

在LSDB中我们看到有2条描述10.1.0.0的路由，接下来查看LSA的详细信息。下面这条第三类LSA是由R2始发的，而这条第五类LSA是由R5始发的。这两条LSA描述了一个完全相同的网段信息。

```
[R5]display ospf lsdb summary 10.1.0.0

       OSPF Process 1 with Router ID 10.0.5.5
                    Area: 0.0.0.0
              Link State Database


  Type      : Sum-Net
  Ls id     : 10.1.0.0
  Adv rtr    : 10.0.2.2
  Ls age     : 767
  Len       : 28
  Options    :  E
  seq#      : 80000001
  chksum     : 0xa380
  Net mask   : 255.255.0.0
  Tos 0  metric: 1562
  Priority   : Low

[R5]display ospf lsdb ase 10.1.0.0

       OSPF Process 1 with Router ID 10.0.5.5
              Link State Database


  Type      : External
```

```
  Ls id     : 10.1.0.0
  Adv rtr   : 10.1.4.4
  Ls age    : 871
  Len       : 36
  Options   :  E
  seq#      : 80000001
  chksum    : 0xe3cd
  Net mask  : 255.255.0.0
  TOS 0  Metric: 2
  E type    : 2
  Forwarding Address : 0.0.0.0
  Tag       : 1
  Priority  : Low
```

在OSPF中，第三类LSA始终优于第五类LSA，所以在R5路由表里出现的10.1.0.0/16这条路由的下一跳会是R3。

为了避免这类问题的发生，我们在R4上取消原来对外部路由的汇总，这样这条路由就会再次在其他路由器的路由表中出现。

```
[R4-ospf-1]undo asbr-summary 10.1.0.0 255.255.0.0

[R5]display ip routing-table protocol os
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Public routing table : OSPF
       Destinations : 6       Routes : 6

OSPF routing table status : <Active>
       Destinations : 6       Routes : 6

Destination/Mask    Proto   Pre  Cost   Flags NextHop       Interface

      10.0.2.0/24  OSPF    10   1563     D   10.0.75.3     GigabitEthernet0/0/0
      10.0.3.0/24  OSPF    10   1        D   10.0.75.3     GigabitEthernet0/0/0
     10.0.12.0/24  OSPF    10   3125     D   10.0.75.3     GigabitEthernet0/0/0
     10.0.23.0/24  OSPF    10   1563     D   10.0.75.3     GigabitEthernet0/0/0
      10.1.0.0/16  OSPF    10   3125     D   10.0.75.3     GigabitEthernet0/0/0
      10.1.4.4/24  O_ASE   150  1        D   10.0.75.4     GigabitEthernet0/0/0

OSPF routing table status : <Inactive>
       Destinations : 0       Routes : 0
```

这时我们可以看到，在R5上已学习到了一条关于10.1.4.4/24正确的路由。

这时我们在R1上测试连通性。

```
[R1]ping -c 1 10.1.4.4
  PING 10.1.4.4: 56  data bytes, press CTRL_C to break
    Reply from 10.1.4.4: bytes=56 Sequence=1 ttl=253 time=71 ms

  --- 10.1.4.4 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 71/71/71 ms
```

可见，网络恢复正常。

附加实验：思考并验证

可否在一个区域中即打开基于区域的认证，又打开基于接口的认证？

非骨干区域的区域号可否一样？

最终设备配置

```
<R1>display current-configuration
[V200R001C00SPC200]
#
 sysname R1
#
interface Serial1/0/0
 link-protocol ppp
 ip address 10.0.12.1 255.255.255.0
 ospf authentication-mode md5 1 plain huawei
#
interface LoopBack0
 ip address 10.1.1.1 255.255.255.0
 ospf network-type broadcast
#
ospf 1 router-id 10.1.1.1
 area 0.0.0.2
  network 10.0.12.1 0.0.0.0
  network 10.1.1.1 0.0.0.0
#
return
```

```
<R2>display current-configuration
[V200R001C00SPC200]
#
 sysname R2
#
interface Serial1/0/0
 link-protocol ppp
 ip address 10.0.12.2 255.255.255.0
 ospf authentication-mode md5 1 plain huawei
#
interface LoopBack0
 ip address 10.0.2.2 255.255.255.0
 ospf network-type broadcast
#
ospf 1 router-id 10.0.2.2
 area 0.0.0.0
  authentication-mode md5 1 plain huawei
 area 0.0.0.1
  network 10.0.23.2 0.0.0.0
  network 10.0.2.2 0.0.0.0
  vlink-peer 10.0.3.3
 area 0.0.0.2
  abr-summary 10.1.0.0 255.255.0.0
  network 10.0.12.2 0.0.0.0
#
return

<R3>display current-configuration
[V200R001C00SPC200]
#
 sysname R3
#
interface Serial2/0/0
 link-protocol ppp
 ip address 10.0.23.3 255.255.255.0
#
interface GigabitEthernet0/0/0
 ip address 10.0.75.3 255.255.255.0
#
interface LoopBack0
 ip address 10.0.3.3 255.255.255.0
 ospf network-type broadcast
#
```

```
ospf 1 router-id 10.0.3.3
 area 0.0.0.0
  authentication-mode md5 1 plain huawei
  network 10.0.75.3 0.0.0.0
 area 0.0.0.1
  network 10.0.23.3 0.0.0.0
  network 10.0.3.3 0.0.0.0
  vlink-peer 10.0.2.2
#
return

<R4>display current-configuration
[V200R001C00SPC200]
#
 sysname R4
#
interface GigabitEthernet0/0/0
 ip address 10.0.75.4 255.255.255.0
#
interface LoopBack0
 ip address 10.1.4.4 255.255.255.0
#
ospf 1 router-id 10.1.4.4
 import-route direct
 area 0.0.0.0
  authentication-mode md5 1 plain huawei
  network 10.0.75.4 0.0.0.0
#
return

<R5>display current-configuration
[V200R001C00SPC200]
#
 sysname R5
#
interface GigabitEthernet0/0/0
 ip address 10.0.75.5 255.255.255.0
#
interface LoopBack0
 ip address 10.0.5.5 255.255.255.0
 ospf network-type broadcast
#
ospf 1 router-id 10.0.5.5
```

```
 area 0.0.0.0
  authentication-mode md5 1 plain huawei
  network 10.0.75.5 0.0.0.0
  network 10.0.5.5 0.0.0.0
#
Return
```

实验 2-7 OSPF 高级特性

学习目的

- 掌握在NBMA网络中手工配置OSPF邻居的方法
- 掌握在NBMA网络中影响DR选举的方法
- 掌握在NBMA网络中配置OSPF的特点
- 掌握在FR中设置广播型网络，使用OSPF的方法
- 掌握在FR中设置P2MP型网络，使用OSPF的方法
- 掌握混合使用P2MP和P2P网络中配置OSPF的方法
- 掌握在FR中建立子接口设置P2P型网络使用OSPF的方法

拓扑图

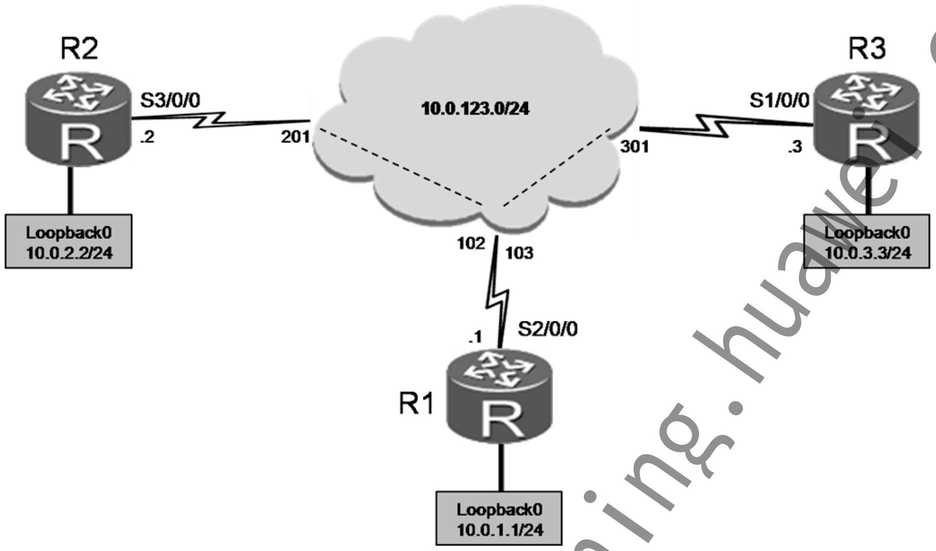

图2-7 OSPF 高级特性

场景

你是公司的网络管理员。公司的网络采用了OSPF协议作为路由协议。公司有三个分支机构，分别有R1、R2、R3三台路由器。出于成本的考虑，在R1和R2之间，R1和R3之间各租用一条虚电路，在R2和R3之间没有虚电路。首先通过配置使OSPF在NBMA网络上运行并观察使用情况，然后再分别将网络修改为广播型、P2MP型，P2MP-P2P混合型以及点对点类型来运行OSPF。

学习任务

步骤一.　基础 FR 互联配置与 IP 编址

给所有路由器配置IP地址和掩码。配置时注意所有的Loopback接口配置掩码均为24位，模拟成一个单独的网段。

默认情况下路由器开启帧中继反向ARP特性（Inverse ARP），需要关闭该特性，并使用手工映射的方法建立R1与R2以及R1与R3之间的ARP映射表。

默认情况下，帧中继线路不能传递广播，为了能在本实验中OSPF邻居发现的正常进行，我们在做帧中继地址映射的时候加上Broadcast参数，允许广播报文通过。

```
<R1>system-view
Enter system view, return user view with Ctrl+Z.
[R1]interface Serial 2/0/0
[R1-Serial2/0/0]link-protocol fr
Warning: The encapsulation protocol of the link will be changed. Continue? [Y/N]:y
[R1-Serial2/0/0]ip address 10.0.123.1 24
[R1-Serial2/0/0]undo fr inarp
[R1-Serial2/0/0]fr map ip 10.0.123.2 102 broadcast
[R1-Serial2/0/0]fr map ip 10.0.123.3 103 broadcast
[R1-Serial2/0/0]interface LoopBack 0
[R1-LoopBack0]ip address 10.0.1.1 24

<R2>system-view
Enter system view, return user view with Ctrl+Z.
[R2]interface Serial 3/0/0
[R2-Serial3/0/0]link-protocol fr
Warning: The encapsulation protocol of the link will be changed. Continue? [Y/N]:y
[R2-Serial3/0/0]ip address 10.0.123.2 24
[R2-Serial3/0/0]undo fr inarp
[R2-Serial3/0/0]fr map ip 10.0.123.1 201 broadcast
[R2-Serial3/0/0]interface LoopBack 0
[R2-LoopBack0]ip address 10.0.2.2 24

<R3>system-view
Enter system view, return user view with Ctrl+Z.
[R3]interface s1/0/0
[R3-Serial1/0/0]link-protocol fr
Warning: The encapsulation protocol of the link will be changed. Continue? [Y/N]:y
[R3-Serial1/0/0]ip address 10.0.123.3 24
[R3-Serial1/0/0]undo fr inarp
[R3-Serial1/0/0]fr map ip 10.0.123.1 301 broadcast
[R3-Serial1/0/0]interface LoopBack 0
[R3-LoopBack0]ip address 10.0.3.3 24
```

配置完成后，**display fr map-info**检查FR的地址映射信息，测试链路的连通性。

```
[R1]display fr map-info
Map Statistics for interface Serial2/0/0 (DTE)
  DLCI = 102, IP 10.0.123.2, Serial2/0/0
    create time = 2011/11/30 09:06:43, status = ACTIVE
    encapsulation = ietf, vlink = 3, broadcast
  DLCI = 103, IP 10.0.123.3, Serial2/0/0
    create time = 2011/11/30 09:06:53, status = ACTIVE
```

```
  encapsulation = ietf, vlink = 4, broadcast

[R1]ping -c 1 10.0.123.2
  PING 10.0.123.2: 56  data bytes, press CTRL_C to break
    Reply from 10.0.123.2: bytes=56 Sequence=1 ttl=255 time=66 ms

  --- 10.0.123.2 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 66/66/66 ms

[R1]ping -c 1 10.0.123.3
  PING 10.0.123.3: 56  data bytes, press CTRL_C to break
    Reply from 10.0.123.3: bytes=56 Sequence=1 ttl=255 time=56 ms

  --- 10.0.123.3 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 56/56/56 ms
```

步骤二.　配置 NBMA 类型的 OSPF 网络

配置10.0.123.0/24及各路由器Loopback 0接口地址属于OSPF区域0。对所有路由器的Loopback0接口，修改其OSPF网络类型为Broadcast类型，使用Loopback 0的地址作为Router ID。

注意在使用network命令时，通配符掩码使用0.0.0.0。

在帧中继网络中，OSPF默认的网络类型为NBMA，对于NBMA网络来说，要求手工配置OSPF的邻居，配置完成以后检查各路由器邻居关系。

```
[R1]ospf 1 router-id 10.0.123.1
[R1-ospf-1]area 0
[R1-ospf-1-area-0.0.0.0]network 10.0.123.1 0.0.0.0
[R1-ospf-1-area-0.0.0.0]network 10.0.1.1 0.0.0.0
[R1-ospf-1-area-0.0.0.0]quit
[R1-ospf-1]peer 10.0.123.2
[R1-ospf-1]peer 10.0.123.3
[R1-ospf-1]interface LoopBack 0
[R1-LoopBack0]ospf network-type broadcast
```

```
[R2]ospf 1 router-id 10.0.2.2
[R2-ospf-1]area 0
[R2-ospf-1-area-0.0.0.0]network 10.0.123.2 0.0.0.0
[R2-ospf-1-area-0.0.0.0]network 10.0.2.2 0.0.0.0
[R2-ospf-1-area-0.0.0.0]quit
[R2-ospf-1]peer 10.0.123.1
[R2-ospf-1]int LoopBack 0
[R2-LoopBack0]ospf network-type broadcast

[R3]ospf 1 router-id 10.0.3.3
[R3-ospf-1]area 0
[R3-ospf-1-area-0.0.0.0]network 10.0.123.3 0.0.0.0
[R3-ospf-1-area-0.0.0.0]network 10.0.3.3 0.0.0.0
[R3-ospf-1-area-0.0.0.0]quit
[R3-ospf-1]peer 10.0.123.1
[R3-ospf-1]interface LoopBack 0
[R3-LoopBack0]ospf network-type broadcast
```

由于R1的OSPF最早配置，所以R1是10.0.123.0这个网段的DR。这时我们通过**reset ospf process**重启R1的OSPF进程。

```
<R1>reset ospf process
Warning: The OSPF process will be reset. Continue? [Y/N]:y

[R2]display ospf peer

        OSPF Process 1 with Router ID 10.0.2.2
               Neighbors

 Area 0.0.0.0 interface 10.0.123.2(Serial3/0/0)'s neighbors
 Router ID: 10.0.123.1          Address: 10.0.123.1
   State: Full  Mode:Nbr is  Master  Priority: 1
   DR: 10.0.123.2  BDR: 10.0.123.1  MTU: 0
   Dead timer due in 33  sec
   Retrans timer interval: 6
   Neighbor is up for 00:01:23
   Authentication Sequence: [ 0 ]
```

这时，看到R2担当了DR的位置。在R2上查看OSPF路由表。

```
[R2]dis ip routing-table protocol ospf
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Public routing table : OSPF
```

```
        Destinations : 1       Routes : 1

OSPF routing table status : <Active>
        Destinations : 1       Routes : 1

Destination/Mask    Proto   Pre  Cost     Flags NextHop        Interface

      10.0.1.0/24  OSPF    10   1562        D   10.0.123.1     Serial3/0/0

OSPF routing table status : <Inactive>
        Destinations : 0       Routes : 0
```

现在可以发现在R2中仅存在一条OSPF路由，为到达R1的Loopback 0连接网段的路由。R2没有学习到R3的Loopback 0所在网段的路由。这是什么原因呢？

NBMA网络本身要求全互联。在本次实验中，R2和R3之间不存在虚电路，没有直接的邻接关系。而这时R2担当DR，所以R3的路由信息是没办法传递给R2。

因此，在这种场景中，我们必须保证R1始终成为DR。我们通过修改R2和R3接口的OSPF优先级来保证R1始终成为DR。

在OSPF中，接口优先级为0的路由器始终不参加DR/BDR的选举。

```
[R2]interface s3/0/0
[R2-Serial3/0/0]ospf dr-priority 0

[R3]interface Serial 1/0/0
[R3-Serial1/0/0]ospf dr-priority 0
```

再次检查OSPF路由表，查看路由是否有缺失；

```
[R2]display ip routing-table protocol ospf
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Public routing table : OSPF
        Destinations : 2       Routes : 2

OSPF routing table status : <Active>
        Destinations : 2       Routes : 2

Destination/Mask    Proto   Pre  Cost     Flags NextHop        Interface

      10.0.1.0/24  OSPF    10   1562        D   10.0.123.1     Serial3/0/0
      10.0.3.0/24  OSPF    10   1562        D   10.0.123.3     Serial3/0/0
```

```
OSPF routing table status : <Inactive>
        Destinations : 0       Routes : 0
```

我们在路由表里已能看到了到达R3的Loopback接口网段的路由，然后在R2上检查连通性。

```
[R2]ping -c 1 10.0.3.3
  PING 10.0.3.3: 56  data bytes, press CTRL_C to break
    Request time out

  --- 10.0.3.3 ping statistics ---
    1 packet(s) transmitted
    0 packet(s) received
    100.00% packet loss
```

我们发现R2无法访问R3的Loopback0接口的地址，这又是什么原因呢？

检查路由表，发现要到达10.0.3.3的路由的下一跳是10.0.123.3。

检查到达该下一跳地址的连通性。

```
[R2]ping -c 1 10.0.123.3
  PING 10.0.123.3: 56  data bytes, press CTRL_C to break
    Request time out

  --- 10.0.123.3 ping statistics ---
    1 packet(s) transmitted
    0 packet(s) received
    100.00% packet loss
```

地址10.0.123.3和R1的接口地址10.0.123.2在同一网段。查看帧中继的映射表。

```
[R2]display fr map-info
Map Statistics for interface Serial3/0/0 (DTE)
  DLCI = 201, IP 10.0.123.1, Serial3/0/0
    create time = 2011/11/30 10:03:37, status = ACTIVE
    encapsulation = ietf, vlink = 1, broadcast
```

发现仅存在到达10.0.123.1的映射，没有到达10.0.123.3的映射关系。

这时我们手工增加R2到R3接口地址的映射，以及R3到R2接口地址的映射。

```
[R2]interface Serial 3/0/0
[R2-Serial3/0/0]fr map ip 10.0.123.3 201
[R2-Serial3/0/0]quit
```

```
[R3]interface Serial 1/0/0
[R3-Serial1/0/0]fr map ip 10.0.123.2 301
[R3-Serial1/0/0]quit
```

再次测试R2与R3之间的连通性。

```
[R2]ping -c 1 10.0.3.3
  PING 10.0.3.3: 56  data bytes, press CTRL_C to break
    Reply from 10.0.3.3: bytes=56 Sequence=1 ttl=254 time=122 ms

  --- 10.0.3.3 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 122/122/122 ms
```

这时R2已能访问R3。

步骤三.　配置广播类型的 OSPF 网络

默认情况下，帧中继线路不能传递广播。为了能在OSPF中使用广播这种网络类型，刚才我们在做帧中继地址映射的时候加上Broadcast参数，允许广播在帧中继线路上通过。

配置OSPF的网络类型为广播，使OSPF在FR接口上，以广播的形式工作。

```
[R1]interface Serial 2/0/0
[R1-Serial2/0/0]ospf network-type broadcast

[R2]interface Serial 3/0/0
[R2-Serial3/0/0]ospf network-type broadcast

[R3]interface Serial 1/0/0
[R3-Serial1/0/0]ospf network-type broadcast
```

在广播型网络中，是不需要手工配置邻居的。所以我们在OSPF进程下删除手工配置的邻居。

```
[R1]ospf 1
[R1-ospf-1]undo peer 10.0.123.2
[R1-ospf-1]undo peer 10.0.123.3

[R2]ospf 1
```

```
[R2-ospf-1]undo peer 10.0.123.1

[R3]ospf 1
[R3-ospf-1]undo peer 10.0.123.1
```

在R1上检查邻居关系，并在R2上查看路由表，测试连通性。

```
[R1]display ospf peer brief

        OSPF Process 1 with Router ID 10.0.123.1
                Peer Statistic Information
 ----------------------------------------------------------------------------
 Area Id          Interface                     Neighbor id     State
 0.0.0.0          Serial2/0/0                    10.0.2.2       Full
 0.0.0.0          Serial2/0/0                    10.0.3.3       Full
 ----------------------------------------------------------------------------

[R2]display ip routing-table protocol ospf
Route Flags: R - relay, D - download to fib
----------------------------------------------------------------------------
Public routing table : OSPF
        Destinations : 2       Routes : 2

OSPF routing table status : <Active>
        Destinations : 2       Routes : 2

Destination/Mask    Proto   Pre  Cost      Flags NextHop        Interface

      10.0.1.0/24  OSPF    10   1562        D   10.0.123.1     Serial3/0/0
      10.0.3.0/24  OSPF    10   1562        D   10.0.123.3     Serial3/0/0

OSPF routing table status : <Inactive>
        Destinations : 0       Routes : 0
```

这时网络已工作正常。

注意，到达网络10.0.3.0/24的下一跳是10.0.123.3，和刚才设置为NBMA的网络类型的结果是一样的。所以对于广播型网络，在R2和R3的接口地址之间还是需要存在帧中继的映射关系的。

在R2上测试网络是否正常。

```
[R2]ping -c 1 10.0.3.3
  PING 10.0.3.3: 56  data bytes, press CTRL_C to break
    Reply from 10.0.3.3: bytes=56 Sequence=1 ttl=254 time=114 ms
```

```
--- 10.0.3.3 ping statistics ---
  1 packet(s) transmitted
  1 packet(s) received
  0.00% packet loss
  round-trip min/avg/max = 114/114/114 ms
```

步骤四. 配置 P2MP 类型的 OSPF 网络

在P2MP这种网络类型中，我们同样不需要手工配置邻居。

把各接口的网络类型修改为P2MP，同时在接口下将设置DR优先级的配置删除。

在P2MP这种类型的网络中，OSPF是不需要进行DR/BDR的选举的。

```
[R1]interface Serial 2/0/0
[R1-Serial2/0/0]ospf network-type p2mp

[R2]interface Serial 3/0/0
[R2-Serial3/0/0]undo ospf dr-priority
[R2-Serial3/0/0]ospf network-type p2mp

[R3]interface Serial 1/0/0
[R3-Serial1/0/0]undo ospf dr-priority
[R3-Serial1/0/0]ospf network-type p2mp
```

在R1上查看邻居关系。

```
[R1]display ospf peer  brief

       OSPF Process 1 with Router ID 10.0.123.1
              Peer Statistic Information
 ----------------------------------------------------------------------------
 Area Id        Interface                      Neighbor id      State
 0.0.0.0        Serial2/0/0                     10.0.2.2        Full
 0.0.0.0        Serial2/0/0                     10.0.3.3        Full
 ----------------------------------------------------------------------------
```

在R1和R2上分别检查路由表，可以看到路由信息均正常传递。

```
[R1]display ip routing-table protocol ospf
Route Flags: R - relay, D - download to fib
----------------------------------------------------------------------------
```

```
Public routing table : OSPF
       Destinations : 4       Routes : 4

OSPF routing table status : <Active>
       Destinations : 2       Routes : 2

Destination/Mask    Proto   Pre  Cost     Flags NextHop        Interface

      10.0.2.0/24  OSPF    10   1562       D   10.0.123.2     Serial2/0/0
      10.0.3.0/24  OSPF    10   1562       D   10.0.123.3     Serial2/0/0

OSPF routing table status : <Inactive>
       Destinations : 2       Routes : 2

Destination/Mask    Proto   Pre  Cost     Flags NextHop        Interface

    10.0.123.2/32  OSPF    10   1562           10.0.123.2     Serial2/0/0
    10.0.123.3/32  OSPF    10   1562           10.0.123.3     Serial2/0/0

[R2]display ip routing-table protocol ospf
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Public routing table : OSPF
       Destinations : 4       Routes : 4

OSPF routing table status : <Active>
       Destinations : 2       Routes : 2

Destination/Mask    Proto   Pre  Cost     Flags NextHop        Interface

      10.0.1.0/24  OSPF    10   1562       D   10.0.123.1     Serial3/0/0
      10.0.3.0/24  OSPF    10   3124       D   10.0.123.1     Serial3/0/0

OSPF routing table status : <Inactive>
       Destinations : 2       Routes : 2

Destination/Mask    Proto   Pre  Cost     Flags NextHop        Interface

    10.0.123.1/32  OSPF    10   1562           10.0.123.1     Serial3/0/0
    10.0.123.3/32  OSPF    10   3124           10.0.123.1     Serial3/0/0
```

观察R2的路由表我们可以发现，到达网段10.0.3.0/24的下一跳变成了10.0.123.1。这样，实际上在R2上只要有对10.0.123.1这个地址帧中继映射就可

以了。

先删除多余的地址映射，再测试连通性。

```
[R2]interface Serial 3/0/0
[R2-Serial3/0/0]undo fr map ip 10.0.123.3 201

[R3]interface Serial 1/0/0
[R3-Serial1/0/0]undo fr map ip 10.0.123.2 301
```

在R2上验证连通性，说明R2到R3的通信正常。

```
[R2]ping -c 1 10.0.3.3
  PING 10.0.3.3: 56  data bytes, press CTRL_C to break
    Reply from 10.0.3.3: bytes=56 Sequence=1 ttl=254 time=119 ms

  --- 10.0.3.3 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 119/119/119 ms
```

在R2上再检查一下删除10.0.123.3映射后路由表有什么变化。

```
[R2-Serial3/0/0]disp ip routing-table protocol ospf
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Public routing table : OSPF
         Destinations : 4        Routes : 4

OSPF routing table status : <Active>
         Destinations : 3        Routes : 3

Destination/Mask    Proto   Pre  Cost      Flags NextHop        Interface

      10.0.1.0/24   OSPF    10   1562        D   10.0.123.1     Serial3/0/0
      10.0.3.0/24   OSPF    10   3124        D   10.0.123.1     Serial3/0/0
    10.0.123.3/32   OSPF    10   3124        D   10.0.123.1     Serial3/0/0

OSPF routing table status : <Inactive>
         Destinations : 1        Routes : 1

Destination/Mask    Proto   Pre  Cost      Flags NextHop        Interface

    10.0.123.1/32  OSPF    10   1562            10.0.123.1     Serial3/0/0
```

这时可以发现刚才10.0.123.3/32不是活跃状态的，因为10.0.123.3/32是以直连路由的形式出现在路由表里的。现在删除帧中继到达10.0.123.3的映射以后，10.0.123.3/32便以OSPF的形式出现在路由表里了。

步骤五. 配置 P2MP 和 P2P 混合的 OSPF 网络

P2MP和P2P两种网络类型可混合存在。

将R2、R3的网络类型更改为P2P，R1的网络类型保持P2MP不变。

```
[R2]interface Serial 3/0/0
[R2-Serial3/0/0]ospf network-type p2p

[R3]interface Serial 1/0/0
[R3-Serial1/0/0]ospf network-type p2p
```

这时我们可以看到路由器之间的邻居关系都消失了，等候一段时间以后也没有再建立。观察OSPF的错误会发现有Hello间隔不匹配的问题出现。

```
Nov 30 2011 14:16:10+00:00 R2 %%01OSPF/3/NBR_CHG_DOWN(l)[0]:Neighbor
event:neighbor state changed to Down. (ProcessId=1, NeighborAddress=10.0.123.1,
NeighborEvent=KillNbr, NeighborPreviousState=Full, NeighborCurrentState=Down)

[R2]display ospf error

        OSPF Process 1 with Router ID 10.0.2.2
               OSPF error statistics

General packet errors:
 0     : IP: received my own packet      6     : Bad packet
 0     : Bad version                     0     : Bad checksum
 0     : Bad area id                     0     : Drop on unnumbered interface
 0     : Bad virtual link                0     : Bad authentication type
 0     : Bad authentication key          0     : Packet too small
 0     : Packet size > ip length         0     : Transmit error
 7     : Interface down                  0     : Unknown neighbor

HELLO packet errors:
 0     : Netmask mismatch                6     : Hello timer mismatch
 0     : Dead timer mismatch             0     : Extern option mismatch
 0     : Router id confusion             0     : Virtual neighbor unknown
 0     : NBMA neighbor unknown           0     : Invalid Source Address
```

对于P2MP网络而言，默认的Hello间隔是30秒。而对于P2P，默认的Hello间隔是10秒。

在这个实验中，调整R1的Hello间隔。

```
[R1]interface Serial 2/0/0
[R1-Serial2/0/0]ospf timer hello 10
```

等待约半分钟后，再次观察邻居关系是否能建立。

```
[R1]display ospf peer brief

        OSPF Process 1 with Router ID 10.0.123.1
                Peer Statistic Information
 ----------------------------------------------------------------------------
 Area Id          Interface                     Neighbor id      State
 0.0.0.0          Serial2/0/0                    10.0.2.2        Full
 0.0.0.0          Serial2/0/0                    10.0.3.3        Full
 ----------------------------------------------------------------------------
```

可以看到路由器之间又形成了邻居关系。

接着检查路由器之间的路由信息传递是否正常，首先在R1上查看路由表。

```
[R1]display ip routing-table protocol ospf
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Public routing table : OSPF
        Destinations : 3        Routes : 4

OSPF routing table status : <Active>
        Destinations : 2        Routes : 2

Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface

      10.0.2.0/24  OSPF    10   1562        D   10.0.123.2      Serial2/0/0
      10.0.3.0/24  OSPF    10   1562        D   10.0.123.3      Serial2/0/0

OSPF routing table status : <Inactive>
        Destinations : 1        Routes : 2

Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface

    10.0.123.0/24  OSPF    10   3124            10.0.123.2      Serial2/0/0
    10.0.123.0/24  OSPF    10   3124            10.0.123.3      Serial2/0/0
```

可以看到R1上有R2和R3的Loopback地址所在网段的路由。

查看R2的路由表。

```
[R2]display ip routing-table protocol ospf
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Public routing table : OSPF
        Destinations : 3       Routes : 3

OSPF routing table status : <Active>
        Destinations : 2       Routes : 2

Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface

      10.0.1.0/24  OSPF    10   1562        D   10.0.123.1      Serial3/0/0
      10.0.3.0/24  OSPF    10   3124        D   10.0.123.1      Serial3/0/0

OSPF routing table status : <Inactive>
        Destinations : 1       Routes : 1

Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface

    10.0.123.1/32  OSPF    10   1562            10.0.123.1      Serial3/0/0
```

在R2上的路由表存在到达R1和R3的路由，但是否能直接访问到呢？

```
[R2]ping -c 1 10.0.3.3
  PING 10.0.3.3: 56  data bytes, press CTRL_C to break
    Request time out

  --- 10.0.3.3 ping statistics ---
    1 packet(s) transmitted
    0 packet(s) received
    100.00% packet loss

[R2]ping -c 1 10.0.1.1
  PING 10.0.1.1: 56  data bytes, press CTRL_C to break
    Reply from 10.0.1.1: bytes=56 Sequence=1 ttl=255 time=63 ms

  --- 10.0.1.1 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 63/63/63 ms
```

这时我们发现R2并不能直接访问到R3，但能直接访问到R1。跟踪到达10.0.3.3的数据包在哪一跳被丢弃。

```
[R2]tracert 10.0.3.3
 traceroute to  10.0.3.3(10.0.3.3), max hops: 30 ,packet length: 40,press CTRL_C
to break
 1 10.0.123.1 61 ms  42 ms  42 ms
 2  *  *  *
 ...
```

上面的结果说明数据包已经到达了R1，但到达R3之后被丢弃了，我们检查R3的全局路由表。

```
[R3]display ip routing-table
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Routing Tables: Public
         Destinations : 13      Routes : 13

Destination/Mask    Proto   Pre  Cost    Flags NextHop         Interface

        10.0.1.0/24    OSPF    10   1562     D  10.0.123.1      Serial1/0/0
        10.0.2.0/24    OSPF    10   3124     D  10.0.123.1      Serial1/0/0
        10.0.3.0/24    Direct  0    0        D  10.0.3.3        LoopBack0
        10.0.3.3/32    Direct  0    0        D  127.0.0.1       InLoopBack0
      10.0.3.255/32    Direct  0    0        D  127.0.0.1       InLoopBack0
      10.0.123.0/24    Direct  0    0        D  10.0.123.3      Serial1/0/0
      10.0.123.1/32    Direct  0    0        D  10.0.123.1      Serial1/0/0
      10.0.123.3/32    Direct  0    0        D  127.0.0.1       InLoopBack0
    10.0.123.255/32    Direct  0    0        D  127.0.0.1       InLoopBack0
       127.0.0.0/8     Direct  0    0        D  127.0.0.1       InLoopBack0
       127.0.0.1/32    Direct  0    0        D  127.0.0.1       InLoopBack0
127.255.255.255/32     Direct  0    0        D  127.0.0.1       InLoopBack0
255.255.255.255/32     Direct  0    0        D  127.0.0.1       InLoopBack0
```

这时我们应该注意到刚才在R2上ping R3的Loopback接口地址的时候，数据包的三层包头的源地址是路由器R2的接口地址，即为10.0.123.2。

在R3的路由表中，没有到达10.0.123.2/32的路由表项，故该数据包被丢弃了。从这里我们可以看出，当网络类型由P2MP改为P2P以后，R2和R3就学习不到对方的直连接口地址了。当然，这其实不影响R2和R3所连接用户的通信，我们可以使用带源地址的方式来验证连通性。

可以看到若使用扩展Ping，将Ping数据包的源地址设为R2的Loopback接口

地址，可以与R3通讯。

```
[R2]ping -c 1 -a 10.0.2.2 10.0.3.3
  PING 10.0.3.3: 56  data bytes, press CTRL_C to break
    Reply from 10.0.3.3: bytes=56 Sequence=1 ttl=254 time=123 ms

  --- 10.0.3.3 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 123/123/123 ms
```

步骤六.　配置 P2P 类型的 OSPF 网络

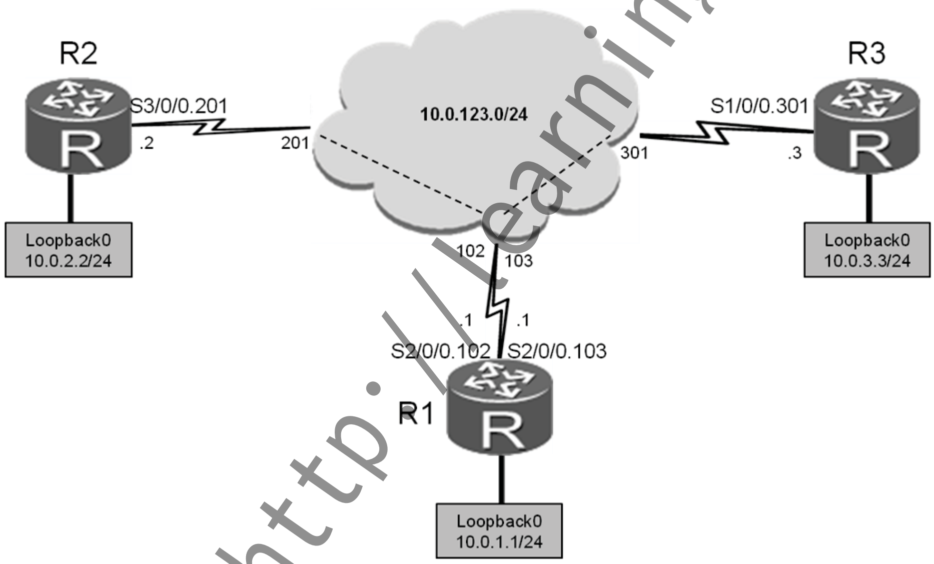

在这里我们要重新对地址做规划。如图所示，R1与R2之间使用10.0.12.0/24互联，R1与R3之间使用网段10.0.13.0/24互联。

配置子接口及地址，在配置子接口时，指定接口类型为P2P。

```
[R1]interface s2/0/0
[R1-Serial2/0/0]undo fr map ip 10.0.123.2 102
[R1-Serial2/0/0]undo fr map ip 10.0.123.3 103
[R1-Serial2/0/0]undo ospf network-type
[R1-Serial2/0/0]undo ospf timer hello
[R1-Serial2/0/0]interface Serial 2/0/0.102  p2p
[R1-Serial2/0/0.102]ip address 10.0.12.1 24
[R1-Serial2/0/0.102]ospf network-type p2p
```

```
[R1-Serial2/0/0.102]fr dlci 102
[R1-fr-dlci-Serial2/0/0.102-102]interface Serial 2/0/0.103 p2p
[R1-Serial2/0/0.103]ip address 10.0.13.1 24
[R1-Serial2/0/0.103]ospf network-type p2p
[R1-Serial2/0/0.103]fr dlci 103

[R2]interface Serial 3/0/0
[R2-Serial3/0/0]undo fr map ip 10.0.123.1 201
[R2-Serial3/0/0]undo ip address
[R2-Serial3/0/0]undo ospf network-type
[R2-Serial3/0/0]interface Serial 3/0/0.201 p2p
[R2-Serial3/0/0.201]ip address 10.0.12.2 24
[R2-Serial3/0/0.201]ospf network-type p2p
[R2-Serial3/0/0.201]fr dlci 201

[R3]interface Serial 1/0/0
[R3-Serial1/0/0]undo ip address
[R3-Serial1/0/0]undo fr map ip 10.0.123.1 301
[R3-Serial1/0/0]undo ospf network-type
[R3-Serial1/0/0]interface Serial 1/0/0.301 p2p
[R3-Serial1/0/0.301]ip address 10.0.13.3 24
[R3-Serial1/0/0.301]ospf network-type p2p
[R3-Serial1/0/0.301]fr dlci 301
```

配置完成以后后检查连通性。

```
[R1]ping -c 1 10.0.12.2
  PING 10.0.12.2: 56  data bytes, press CTRL_C to break
    Reply from 10.0.12.2: bytes=56 Sequence=1 ttl=255 time=59 ms

  --- 10.0.12.2 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 59/59/59 ms

[R1]ping -c 1 10.0.13.3
  PING 10.0.13.3: 56  data bytes, press CTRL_C to break
    Reply from 10.0.13.3: bytes=56 Sequence=1 ttl=255 time=59 ms

  --- 10.0.13.3 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
```

```
round-trip min/avg/max = 59/59/59 ms
```

在OSPF中修改需发布的网段信息。

```
[R1]ospf 1
[R1-ospf-1]area 0
[R1-ospf-1-area-0.0.0.0]undo network 10.0.123.1 0.0.0.0
[R1-ospf-1-area-0.0.0.0]network 10.0.12.1 0.0.0.0
[R1-ospf-1-area-0.0.0.0]network 10.0.13.1 0.0.0.0

[R2]ospf 1
[R2-ospf-1]area 0
[R2-ospf-1-area-0.0.0.0]undo network 10.0.123.2 0.0.0.0
[R2-ospf-1-area-0.0.0.0]network 10.0.12.2 0.0.0.0

[R3]ospf 1
[R3-ospf-1]area 0
[R3-ospf-1-area-0.0.0.0]undo  network 10.0.123.3 0.0.0.0
[R3-ospf-1-area-0.0.0.0]network  10.0.13.3 0.0.0.0
```

检查OSPF邻居表。

```
[R1]display ospf peer brief

       OSPF Process 1 with Router ID 10.0.123.1
              Peer Statistic Information
 ----------------------------------------------------------------------------
 Area Id          Interface                      Neighbor id      State
 0.0.0.0          Serial2/0/0.102                 10.0.2.2         Full
 0.0.0.0          Serial2/0/0.103                 10.0.3.3         Full
 ----------------------------------------------------------------------------
```

然后在R1和R2上分别检查OSPF路由表。

```
[R1]display ip routing-table protocol ospf
Route Flags: R - relay, D - download to fib
----------------------------------------------------------------------------
Public routing table : OSPF
       Destinations : 2       Routes : 2

OSPF routing table status : <Active>
       Destinations : 2       Routes : 2

Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface
```

```
      10.0.2.0/24  OSPF    10   1562       D   10.0.12.2      Serial2/0/0.102
      10.0.3.0/24  OSPF    10   1562       D   10.0.13.3      Serial2/0/0.103

OSPF routing table status : <Inactive>
       Destinations : 0       Routes : 0

[R2]display ip routing-table protocol ospf
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Public routing table : OSPF
       Destinations : 3       Routes : 3

OSPF routing table status : <Active>
       Destinations : 3       Routes : 3

Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface

      10.0.1.0/24  OSPF    10   1562       D   10.0.12.1      Serial3/0/0.201
      10.0.3.0/24  OSPF    10   3124       D   10.0.12.1      Serial3/0/0.201
     10.0.13.0/24  OSPF    10   3124       D   10.0.12.1      Serial3/0/0.201

OSPF routing table status : <Inactive>
       Destinations : 0       Routes : 0
```

最后检查网络的连通性。

```
[R2]ping -c 1 10.0.1.1
  PING 10.0.1.1: 56  data bytes, press CTRL_C to break
    Reply from 10.0.1.1: bytes=56 Sequence=1 ttl=255 time=65 ms

  --- 10.0.1.1 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 65/65/65 ms

[R2]ping -c 1 10.0.3.3
  PING 10.0.3.3: 56  data bytes, press CTRL_C to break
    Reply from 10.0.3.3: bytes=56 Sequence=1 ttl=255 time=95 ms

  --- 10.0.3.3 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
```

```
0.00% packet loss
round-trip min/avg/max = 95/95/95 ms
```

附加实验：思考并验证

NBMA网络在非全互联的网络上使用时，有什么注意事项？

比较NBMA、P2MP和P2P在使用上的区别。

最终设备配置

```
<R1>display current-configuration
[V200R001C00SPC200]
#
 sysname R1
#
interface Serial2/0/0
 link-protocol fr
 undo fr inarp
#
interface Serial2/0/0.102 p2p
 fr dlci 102
 ip address 10.0.12.1 255.255.255.0
 ospf network-type p2p
#
interface Serial2/0/0.103 p2p
 fr dlci 103
 ip address 10.0.13.1 255.255.255.0
 ospf network-type p2p
#
interface LoopBack0
 ip address 10.0.1.1 255.255.255.0
 ospf network-type broadcast
#
ospf 1 router-id 10.0.123.1
 area 0.0.0.0
  network 10.0.1.1 0.0.0.0
  network 10.0.12.1 0.0.0.0
  network 10.0.13.1 0.0.0.0
#
return
```

```
<R2>display current-configuration
[V200R001C00SPC200]
#
 sysname R2
#
interface Serial3/0/0
 link-protocol fr
 undo fr inarp
#
interface Serial3/0/0.201 p2p
 fr dlci 201
 ip address 10.0.12.2 255.255.255.0
 ospf network-type p2p
#
interface LoopBack0
 ip address 10.0.2.2 255.255.255.0
 ospf network-type broadcast
#
ospf 1 router-id 10.0.2.2
 area 0.0.0.0
  network 10.0.2.2 0.0.0.0
  network 10.0.12.2 0.0.0.0
#
return

<R3>display current-configuration
[V200R001C00SPC200]
#
 sysname R3
#
interface Serial1/0/0
 link-protocol fr
 fr map ip 10.0.13.1 301 broadcast
 ip address 10.0.13.3 255.255.255.0
 ospf network-type p2p
#
interface LoopBack0
 ip address 10.0.3.3 255.255.255.0
 ospf network-type broadcast
#
ospf 1 router-id 10.0.3.3
 area 0.0.0.0
  network 10.0.3.3 0.0.0.0
```

```
 network 10.0.13.3 0.0.0.0
#
return
```

第三章 BGP协议特性与配置

实验 3-1 IBGP 与 EBGP

学习目的

- 掌握区域内部BGP的配置方法
- 掌握多区域BGP的配置方法
- 观察BGP的邻居表和数据库
- 掌握BGP更新源的配置方法
- 掌握EBGP多跳的配置方法
- 观察IBGP和EBGP中路由的下一跳的变化
- 掌握IBGP中下一跳的配置
- 掌握BGP的Network命令的配置方法

拓扑图

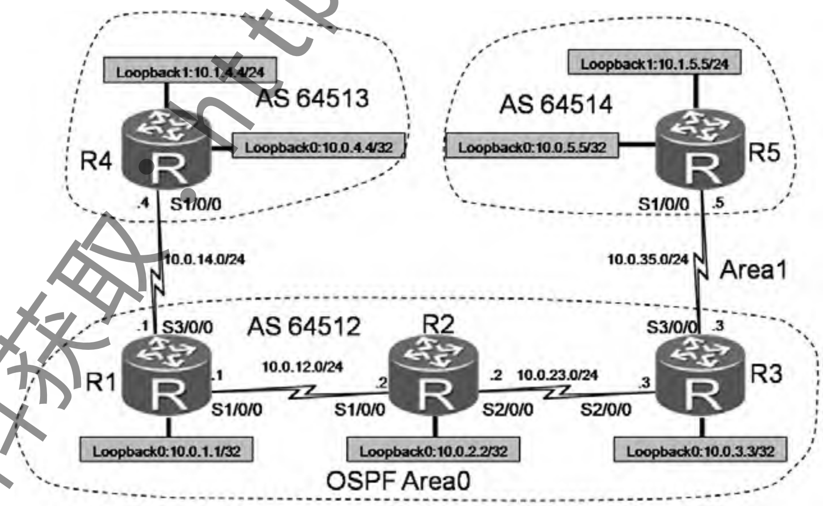

图3-1 OSPF 故障排除

场景

你是公司的网络管理员。公司的网络采用了BGP协议作为路由协议。公司的网络由多个自治系统组成，不同的分支机构使用了不同的AS号，现在你需要完成公司网络的搭建工作。在公司总部使用了OSPF作为IGP，公司内部不同分支机构使用的是私有的BGP AS号。在完成网络搭建以后，你还需要观察BGP路由信息的传递。

学习任务

步骤一. 基础配置与 IP 编址

给所有路由器配置IP地址和掩码，其中R4和R5的loopback 1接口掩码为24位，模拟用户网络。

```
<R1>system-view
Enter system view, return user view with Ctrl+Z.
[R1]interface Serial 1/0/0
[R1-Serial1/0/0]ip address 10.0.12.1 24
[R1-Serial1/0/0]interface Serial 3/0/0
[R1-Serial3/0/0]ip address 10.0.14.1 24
[R1-Serial3/0/0]interface LoopBack 0
[R1-LoopBack0]ip address 10.0.1.1 32

<R2>system-view
Enter system view, return user view with Ctrl+Z.
[R2]interface Serial 1/0/0
[R2-Serial1/0/0]ip address 10.0.12.2 24
[R2-Serial1/0/0]interface Serial 2/0/0
[R2-Serial2/0/0]ip address 10.0.23.2 24
[R2-Serial2/0/0]interface LoopBack 0
[R2-LoopBack0]ip address 10.0.2.2 32

<R3>system-view
Enter system view, return user view with Ctrl+Z.
[R3]interface Serial 2/0/0
[R3-Serial2/0/0]ip address 10.0.23.3 24
[R3-Serial2/0/0]interface Serial 3/0/0
[R3-Serial3/0/0]ip address 10.0.35.3 24
[R3-Serial3/0/0]interface LoopBack 0
[R3-LoopBack0]ip address 10.0.3.3 32
```

```
<R4>system-view
Enter system view, return user view with Ctrl+Z.
[R4]interface Serial 1/0/0
[R4-Serial1/0/0]ip address 10.0.14.4 24
[R4-Serial1/0/0]interface LoopBack 0
[R4-LoopBack0]ip address 10.0.4.4 32

<R5>system-view
Enter system view, return user view with Ctrl+Z.
[R5]interface Serial 1/0/0
[R5-Serial1/0/0]ip address 10.0.35.5 24
[R5-Serial1/0/0]interface LoopBack 0
[R5-LoopBack0]ip address 10.0.5.5 32
```

测试各直连链路的连通性。

```
<R1>ping -c 1 10.0.12.2
  PING 10.0.12.2: 56  data bytes, press CTRL_C to break
    Reply from 10.0.12.2: bytes=56 Sequence=1 ttl=255 time=34 ms

  --- 10.0.12.2 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 34/34/34 ms

<R1>ping -c 1 10.0.14.4
  PING 10.0.14.4: 56  data bytes, press CTRL_C to break
    Reply from 10.0.14.4: bytes=56 Sequence=1 ttl=255 time=40 ms

  --- 10.0.14.4 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 40/40/40 ms

<R3>ping -c 1 10.0.23.2
  PING 10.0.23.2: 56  data bytes, press CTRL_C to break
    Reply from 10.0.23.2: bytes=56 Sequence=1 ttl=255 time=33 ms

  --- 10.0.23.2 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
```

```
    0.00% packet loss
    round-trip min/avg/max = 33/33/33 ms

<R3>ping -c 1 10.0.35.5
  PING 10.0.35.5: 56  data bytes, press CTRL_C to break
    Reply from 10.0.35.5: bytes=56 Sequence=1 ttl=255 time=35 ms

  --- 10.0.35.5 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 35/35/35 ms
```

显示直连联络连通性正常。

步骤二. 配置区域内 IGP

在AS 64512中使用OSPF作为IGP，将Loopback 0连接的网段发布进OSPF。R1的S1/0/0连接的网段运行OSPF。

```
[R1]router id 10.0.1.1
[R1]ospf 1
[R1-ospf-1]area 0
[R1-ospf-1-area-0.0.0.0]network 10.0.12.1 0.0.0.0
[R1-ospf-1-area-0.0.0.0]network 10.0.1.1 0.0.0.0
```

R2的S1/0/0和S2/0/0连接的网段运行OSPF。

```
[R2]router id 10.0.2.2
[R2]ospf 1
[R2-ospf-1]area 0
[R2-ospf-1-area-0.0.0.0]network 10.0.12.2 0.0.0.0
[R2-ospf-1-area-0.0.0.0]network 10.0.23.2  0.0.0.0
[R2-ospf-1-area-0.0.0.0]network 10.0.2.2 0.0.0.0
```

R3的S2/0/0连接的网段运行OSPF。

```
[R3]router id 10.0.3.3
[R3]ospf 1
[R3-ospf-1]area 0
[R3-ospf-1-area-0.0.0.0]network 10.0.23.3 0.0.0.0
[R3-ospf-1-area-0.0.0.0]network 10.0.3.3 0.0.0.0
```

注意在使用network命令时，通配符掩码使用0.0.0.0。

配置完成后察看OSPF的邻居关系是否建立。

```
[R2]display ospf peer

         OSPF Process 1 with Router ID 10.0.2.2
                 Neighbors

 Area 0.0.0.0 interface 10.0.12.2(Serial1/0/0)'s neighbors
 Router ID: 10.0.1.1         Address: 10.0.12.1
   State: Full  Mode:Nbr is  Slave  Priority: 1
   DR: None   BDR: None   MTU: 0
   Dead timer due in 31  sec
   Retrans timer interval: 4
   Neighbor is up for 00:00:29
   Authentication Sequence: [ 0 ]

                 Neighbors

 Area 0.0.0.0 interface 10.0.23.2(Serial2/0/0)'s neighbors
 Router ID: 10.0.3.3         Address: 10.0.23.3
   State: Full  Mode:Nbr is  Master  Priority: 1
   DR: None   BDR: None   MTU: 0
   Dead timer due in 34  sec
   Retrans timer interval: 4
   Neighbor is up for 00:00:06
   Authentication Sequence: [ 0 ]
```

查看所有路由器的路由表。检查是否学习到对端设备Loopback接口网段的路由。

```
[R1]display ip routing-table
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Routing Tables: Public
         Destinations : 15      Routes : 15

Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface

       10.0.1.1/32  Direct  0    0           D   127.0.0.1       InLoopBack0
       10.0.2.2/32  OSPF    10   1562        D   10.0.12.2       Serial1/0/0
       10.0.3.3/32  OSPF    10   3124        D   10.0.12.2       Serial1/0/0
      10.0.12.0/24  Direct  0    0           D   10.0.12.1       Serial1/0/0
      10.0.12.1/32  Direct  0    0           D   127.0.0.1       InLoopBack0
      10.0.12.2/32  Direct  0    0           D   10.0.12.2       Serial1/0/0
```

```
   10.0.12.255/32  Direct  0    0           D   127.0.0.1       InLoopBack0
     10.0.14.0/24  Direct  0    0           D   10.0.14.1       Serial3/0/0
     10.0.14.1/32  Direct  0    0           D   127.0.0.1       InLoopBack0
     10.0.14.4/32  Direct  0    0           D   10.0.14.4       Serial3/0/0
   10.0.14.255/32  Direct  0    0           D   127.0.0.1       InLoopBack0
     10.0.23.0/24  OSPF    10   3124        D   10.0.12.2       Serial1/0/0
     127.0.0.0/8   Direct  0    0           D   127.0.0.1       InLoopBack0
     127.0.0.1/32  Direct  0    0         D   127.0.0.1       InLoopBack0
127.255.255.255/32 Direct  0    0         D   127.0.0.1       InLoopBack0
255.255.255.255/32 Direct  0    0         D   127.0.0.1       InLoopBack0

[R2]display ip routing-table
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Routing Tables: Public
        Destinations : 15       Routes : 15

Destination/Mask     Proto   Pre  Cost      Flags NextHop         Interface

     10.0.1.1/32     OSPF    10   1562        D   10.0.12.1       Serial1/0/0
     10.0.2.2/32     Direct  0    0           D   127.0.0.1       InLoopBack0
     10.0.3.3/32     OSPF    10   1562        D   10.0.23.3       Serial2/0/0
     10.0.12.0/24    Direct  0    0           D   10.0.12.2       Serial1/0/0
     10.0.12.1/32    Direct  0    0           D   10.0.12.1       Serial1/0/0
     10.0.12.2/32    Direct  0    0           D   127.0.0.1       InLoopBack0
     10.0.12.255/32  Direct  0    0           D   127.0.0.1       InLoopBack0
     10.0.23.0/24    Direct  0    0           D   10.0.23.2       Serial2/0/0
     10.0.23.2/32    Direct  0    0           D   127.0.0.1       InLoopBack0
     10.0.23.3/32    Direct  0    0           D   10.0.23.3       Serial2/0/0
     10.0.23.255/32  Direct  0    0           D   127.0.0.1       InLoopBack0
     127.0.0.0/8     Direct  0    0           D   127.0.0.1       InLoopBack0
     127.0.0.1/32    Direct  0    0           D   127.0.0.1       InLoopBack0
127.255.255.255/32   Direct  0    0           D   127.0.0.1       InLoopBack0
255.255.255.255/32   Direct  0    0           D   127.0.0.1       InLoopBack0

[R3]display ip routing-table
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Routing Tables: Public
        Destinations : 16       Routes : 16

Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface
```

```
      10.0.1.1/32   OSPF    10   3124        D   10.0.23.2        Serial2/0/0
      10.0.2.2/32   OSPF    10   1562        D   10.0.23.2        Serial2/0/0
      10.0.3.3/32   Direct  0    0           D   127.0.0.1        InLoopBack0
     10.0.12.0/24   OSPF    10   3124        D   10.0.23.2        Serial2/0/0
     10.0.23.0/24   Direct  0    0           D   10.0.23.3        Serial2/0/0
     10.0.23.2/32   Direct  0    0           D   10.0.23.2        Serial2/0/0
     10.0.23.3/32   Direct  0    0           D   127.0.0.1        InLoopBack0
   10.0.23.255/32   Direct  0    0           D   127.0.0.1        InLoopBack0
     10.0.35.0/24   Direct  0    0           D   10.0.35.3        Serial3/0/0
     10.0.35.3/32   Direct  0    0           D   127.0.0.1        InLoopBack0
     10.0.35.5/32   Direct  0    0           D   10.0.35.5        Serial3/0/0
   10.0.35.255/32   Direct  0    0           D   127.0.0.1        InLoopBack0
     127.0.0.0/8    Direct  0    0           D   127.0.0.1        InLoopBack0
     127.0.0.1/32   Direct  0    0           D   127.0.0.1        InLoopBack0
127.255.255.255/32  Direct  0    0           D   127.0.0.1        InLoopBack0
255.255.255.255/32  Direct  0    0           D   127.0.0.1        InLoopBack0
```

从R1，R2，R3的路由表中，可以看到学习每一台路由器都能够学到其他2台路由器的loopback 0接口连接的网段的路由。

步骤三. 建立 IBGP 对等体

在R1、R2、R3上配置IBGP全互联。使用Loopback0地址作为更新源。

```
[R1]bgp 64512
[R1-bgp]peer 10.0.2.2 as-number 64512
[R1-bgp]peer 10.0.2.2 connect-interface LoopBack 0
[R1-bgp]peer 10.0.3.3 as-number 64512
[R1-bgp]peer 10.0.3.3 connect-interface LoopBack 0

[R2]bgp 64512
[R2-bgp]peer 10.0.1.1 as-number 64512
[R2-bgp]peer 10.0.1.1 connect-interface loopback 0
[R2-bgp]peer 10.0.3.3 as-number 64512
[R2-bgp]peer 10.0.3.3 connect-interface LoopBack 0

[R3]bgp 64512
[R3-bgp]peer 10.0.1.1 as-number 64512
[R3-bgp]peer 10.0.1.1 connect-interface loopback 0
[R3-bgp]peer 10.0.2.2 as-number 64512
[R3-bgp]peer 10.0.2.2 connect-interface LoopBack 0
```

使用**display tcp status**查看TCP端口连接状态。

```
[R2]display tcp status
TCPCB     Tid/Soid Local Add:port       Foreign Add:por   VPNID  State
194a3c7c  8  /2     0.0.0.0:22           0.0.0.0:0        23553   Listening
194a3b18  8  /1     0.0.0.0:23           0.0.0.0:0        23553   Listening
194a3850  106/1     0.0.0.0:80           0.0.0.0:0            0   Listening
19ec2bb8  234/2     0.0.0.0:179          10.0.1.1:0           0   Listening
19ec2360  234/5     0.0.0.0:179          10.0.3.3:0           0   Listening
194a3de0  8  /3     0.0.0.0:830          0.0.0.0:0        23553   Listening
194a39b4  6  /1     0.0.0.0:7547         0.0.0.0:0            0   Listening
19ec3410  234/11    10.0.2.2:179         10.0.3.3:49663       0   Established
19ec2a54  234/4     10.0.2.2:50151       10.0.1.1:179         0   Established
```

从表项中我们可以观察到Local Add为10.0.2.2（即为R2的Loopback0接口地址），端口号为179（BGP协议的TCP端口号）。与10.0.3.3和10.0.1.1的状态已经为Established，说明R2和R1，R3的TCP连接已建立。

使用**display bgp peer**察看各路由器BGP邻居关系状态。

```
[R1]display bgp peer

 BGP local router ID : 10.0.1.1
 Local AS number : 64512
 Total number of peers : 2            Peers in established state : 2

  Peer            V          AS  MsgRcvd  MsgSent  OutQ  Up/Down     State  PrefRcv

  10.0.2.2        4       64512      273      277     0 02:15:53 Established       0
  10.0.3.3        4       64512      276      276     0 02:15:53 Established       0


[R2]display bgp peer

 BGP local router ID : 10.0.2.2
 Local AS number : 64512
 Total number of peers : 2            Peers in established state : 2

  Peer            V          AS  MsgRcvd  MsgSent  OutQ  Up/Down    State    PrefRcv

  10.0.1.1        4       64512       38       38    0 00:18:02 Established       0
  10.0.3.3        4       64512     1000     1000     0 16:38:38 Established       0

[R3]display bgp peer
```

```
BGP local router ID : 10.0.3.3
Local AS number : 64512
Total number of peers : 2          Peers in established state : 2

 Peer            V          AS  MsgRcvd  MsgSent  OutQ  Up/Down     State  PrefRcv

 10.0.1.1        4       64512       39       39     0 00:18:35 Established      0
 10.0.2.2        4       64512     1001     1001     0 16:39:11 Established      0
```

可以看到3台路由器之间的BGP邻居关系处于Established状态，说明邻居关系已建立。

在R1上在BGP进程下使用**timer**修改BGP的keep alive时间为30秒，hold时间为90秒。观察R1与R2的对等体关系建立是否有问题，使用**display bgp peer verbose**命令观察建立以后协商的间隔时间是多少。

```
[R1-bgp] timer keepalive 30 hold 90
```

注意：修改此参数将引起bgp邻居重启。

```
[R2]display bgp peer verbose

        BGP Peer is 10.0.1.1,  remote AS 64512
        Type: IBGP link
        BGP version 4, Remote router ID 10.0.1.1
        Update-group ID: 1
        BGP current state: Established, Up for 00h07m19s
        BGP current event: KATimerExpired
        BGP last state: OpenConfirm
        BGP Peer Up count: 2
        Received total routes: 0
        Received active routes total: 0
        Advertised total routes: 0
        Port:  Local - 50117    Remote - 179
        Configured: Connect-retry Time: 32 sec
        Configured: Active Hold Time: 180 sec   Keepalive Time:60 sec
        Received  : Active Hold Time: 90 sec
        Negotiated: Active Hold Time: 90 sec   Keepalive Time:30 sec
        Peer optional capabilities:
        Peer supports bgp multi-protocol extension
        Peer supports bgp route refresh capability
        Peer supports bgp 4-byte-as capability
        Address family IPv4 Unicast: advertised and received
 Received: Total 16 messages
```

```
            Update messages             0
            Open messages              1
            KeepAlive messages          15
            Notification messages        0
            Refresh messages            0
Sent: Total 16 messages
            Update messages             0
            Open messages              1
            KeepAlive messages          15
            Notification messages        0
            Refresh messages            0
Authentication type configured: None
Last keepalive received: 2011/12/07 08:33:52
Minimum route advertisement interval is 15 seconds
Optional capabilities:
Route refresh capability has been enabled
4-byte-as capability has been enabled
Connect-interface has been configured
Peer Preferred Value: 0
Routing policy configured:
No routing policy is configured

      BGP Peer is 10.0.3.3,  remote AS 64512
      Type: IBGP link
      BGP version 4, Remote router ID 10.0.3.3
      Update-group ID: 1
      BGP current state: Established, Up for 16h28m14s
      BGP current event: RecvKeepalive
      BGP last state: OpenConfirm
      BGP Peer Up count: 1
      Received total routes: 0
      Received active routes total: 0
      Advertised total routes: 0
      Port:  Local - 179     Remote - 49663
      Configured: Connect-retry Time: 32 sec
      Configured: Active Hold Time: 180 sec   Keepalive Time:60 sec
      Received  : Active Hold Time: 180 sec
      Negotiated: Active Hold Time: 180 sec   Keepalive Time:60 sec
      Peer optional capabilities:
      Peer supports bgp multi-protocol extension
      Peer supports bgp route refresh capability
      Peer supports bgp 4-byte-as capability
      Address family IPv4 Unicast: advertised and received
```

```
Received: Total 990 messages
             Update messages            0
             Open messages             1
             KeepAlive messages          989
             Notification messages        0
             Refresh messages           0
Sent: Total 990 messages
             Update messages            0
             Open messages             1
             KeepAlive messages          989
             Notification messages        0
             Refresh messages           0
Authentication type configured: None
Last keepalive received: 2011/12/07 08:34:17
Minimum route advertisement interval is 15 seconds
Optional capabilities:
Route refresh capability has been enabled
4-byte-as capability has been enabled
Connect-interface has been configured
Peer Preferred Value: 0
Routing policy configured:
No routing policy is configured
```

可以看到在R2上默认的配置参数Active Hold Time为180s，Keepalive Time为60s。

当R1的参数修改之后，R2收到数据包的Active Hold Time为90s。协商的参数取值数值小的参数，所以R2与R1的协商的结果Active Hold Time为90s，keepalive Time为30s，而R3的参数仍然为默认参数。

R2和R3一样，所以协商的结果配置参数和协商参数一致，Active Hold Time为180s，Keepalive Time为60s。

步骤四.　配置 EBGP 对等体

在R4上配置BGP，本地AS号为64513，与R1建立EBGP对等体关系。在建立对等体关系时，指定更新源为Loopback 0接口的地址，并指定**ebgp-max-hop**为2。添加到对端Loopback 0接口地址的32位的静态路由，使之能正常建立对等体关系。

```
[R1]ip route-static 10.0.4.4 32 10.0.14.4
[R4]ip route-static 10.0.1.1 32 10.0.14.1

[R1]bgp 64512
```

```
[R1-bgp]peer 10.0.4.4 as-number 64513
[R1-bgp]peer 10.0.4.4 ebgp-max-hop 2
[R1-bgp]peer 10.0.4.4 connect-interface LoopBack0

[R4]bgp 64513
[R4-bgp]peer 10.0.1.1 as-number 64512
[R4-bgp]peer 10.0.1.1 ebgp-max-hop 2
[R4-bgp]peer 10.0.1.1 connect-interface LoopBack0
```

对等体关系建立完成后，使用**display bgp peer**检查对等体关系状态。

```
[R4]display bgp peer

 BGP local router ID : 10.0.4.4
 Local AS number : 64513
 Total number of peers : 1              Peers in established state : 1

Peer         V       AS  MsgRcvd   MsgSent    OutQ    Up/Down  State PrefRcv

10.0.1.1      4       64512      4        5       0 00:01:18    Established       0
```

在R4上使用**debugging ip packet verbose**查看keepalive报文的TTL值。

```
<R4>debugging ip packet verbose
Dec  7 2011 09:09:07.240.2+00:00 R4 IP/7/debug_case:
Delivering, interface = S1/0/0, version = 4, headlen = 20, tos = 192,
pktlen = 40, pktid = 11346, offset = 0, ttl = 2, protocol = 6,
checksum = 29370, s = 10.0.1.1, d = 10.0.4.4
prompt: Packet is before IP_Reass before really deliver to up.

45 c0 00 28 2c 52 00 00 02 06 72 ba 0a 00 01 01
0a 00 04 04 c7 cd 00 b3 91 99 51 7b 2b aa b0 8f
50 10 40 00 cf 00 00 00

Dec  7 2011 09:11:07.640.3+00:00 R4 IP/7/debug_case:
Delivering, interface = S1/0/0, version = 4, headlen = 20, tos = 192,
pktlen = 40, pktid = 11383, offset = 0, ttl = 2, protocol = 6,
checksum = 29333, s = 10.0.1.1, d = 10.0.4.4
prompt: IP packet is delivering up!
```

可以看到收到的报文中TTL都为2。

在R3和R5之间也建立EBGP对等体关系。直接使用物理接口地址建立连接。

```
[R3]bgp 64512
```

```
[R3-bgp]peer 10.0.35.5 as-number 64514

[R5]bgp 64514
[R5-bgp]peer 10.0.35.3 as-number 64512
[R5-bgp]display bgp peer

 BGP local router ID : 10.1.5.5
   Local AS number : 64514
   Total number of peers : 1              Peers in established state : 1

   Peer        V     AS     MsgRcvd  MsgSent   OutQ  Up/Down    State     PrefRcv

    10.0.35.3    4    64512      2      2        0 00:00:57   Established       0
```

步骤五. 使用 Network 命令发布路由信息

在R4上配置Loopback1，地址为10.1.4.4/24。使用**network**命令将该网段发布进BGP。

```
[R4]interface LoopBack 1
[R4-LoopBack1]ip address 10.1.4.4 24
[R4-LoopBack1]bgp 64513
[R4-bgp]network 10.1.4.4 24
```

在R1和R3上全局路由表分别查看该路由是否存在。

查看R3上BGP路由表，分析该路由的下一跳信息。

```
[R1]display ip routing-table
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Routing Tables: Public
        Destinations : 18      Routes : 18

Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface

      10.0.1.1/32   Direct  0    0           D   127.0.0.1       InLoopBack0
      10.0.2.2/32   OSPF    10   1562        D   10.0.12.2       Serial1/0/0
      10.0.3.3/32   OSPF    10   3124        D   10.0.12.2       Serial1/0/0
      10.0.4.4/32   Static  60   0          RD   10.0.14.4       Serial3/0/0
     10.0.12.0/24   Direct  0    0           D   10.0.12.1       Serial1/0/0
     10.0.12.1/32   Direct  0    0           D   127.0.0.1       InLoopBack0
     10.0.12.2/32   Direct  0    0           D   10.0.12.2       Serial1/0/0
```

```
   10.0.12.255/32  Direct  0    0          D   127.0.0.1       InLoopBack0
     10.0.14.0/24  Direct  0    0          D   10.0.14.1       Serial3/0/0
     10.0.14.1/32  Direct  0    0          D   127.0.0.1       InLoopBack0
     10.0.14.4/32  Direct  0    0          D   10.0.14.4       Serial3/0/0
   10.0.14.255/32  Direct  0    0          D   127.0.0.1       InLoopBack0
     10.0.23.0/24  OSPF    10   3124       D   10.0.12.2       Serial1/0/0
      10.1.4.0/24  EBGP    255  0          RD  10.0.4.4        Serial3/0/0
      127.0.0.0/8  Direct  0    0          D   127.0.0.1       InLoopBack0
     127.0.0.1/32  Direct  0    0          D   127.0.0.1       InLoopBack0
127.255.255.255/32 Direct  0    0          D   127.0.0.1       InLoopBack0
255.255.255.255/3  Direct  0    0          D   127.0.0.1       InLoopBack0
```

可以看到在R1上已经学到10.1.4.0/24的EBGP路由。

在R3上查看是否有到达网络10.1.4.0/24的路由。

```
[R3]display ip routing-table
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Routing Tables: Public
         Destinations : 16      Routes : 16

Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface

      10.0.1.1/32  OSPF    10   3124       D   10.0.23.2       Serial2/0/0
      10.0.2.2/32  OSPF    10   1562       D   10.0.23.2       Serial2/0/0
      10.0.3.3/32  Direct  0    0          D   127.0.0.1       InLoopBack0
     10.0.12.0/24  OSPF    10   3124       D   10.0.23.2       Serial2/0/0
     10.0.23.0/24  Direct  0    0          D   10.0.23.3       Serial2/0/0
     10.0.23.2/32  Direct  0    0          D   10.0.23.2       Serial2/0/0
     10.0.23.3/32  Direct  0    0          D   127.0.0.1       InLoopBack0
   10.0.23.255/32  Direct  0    0          D   127.0.0.1       InLoopBack0
     10.0.35.0/24  Direct  0    0          D   10.0.35.3       Serial3/0/0
     10.0.35.3/32  Direct  0    0          D   127.0.0.1       InLoopBack0
     10.0.35.5/32  Direct  0    0          D   10.0.35.5       Serial3/0/0
   10.0.35.255/32  Direct  0    0          D   127.0.0.1       InLoopBack0
      127.0.0.0/8  Direct  0    0          D   127.0.0.1       InLoopBack0
     127.0.0.1/32  Direct  0    0          D   127.0.0.1       InLoopBack0
127.255.255.255/32 Direct  0    0          D   127.0.0.1       InLoopBack0
255.255.255.255/32 Direct  0    0          D   127.0.0.1       InLoopBack0
```

在R3上并没有10.1.4.4的bgp路由。

查看R3的BGP表。

```
[R3]display bgp routing-table

 BGP Local router ID is 10.0.3.3
 Status codes: * - valid, > - best, d - damped,
               h - history,  i - internal, s - suppressed, S - Stale
               Origin : i - IGP, e - EGP, ? - incomplete



 Total Number of Routes: 1
     Network          NextHop        MED        LocPrf    PrefVal Path/Ogn

   i  10.1.4.0/24       10.0.4.4       0          100        0      64513i
```

可以在R3的BGP路由表中看到，但是这条BGP路由没有*号标识，说明这条路由并没有被优选。因为这条路由的NextHop为10.0.4.4，而R3上并没有到达地址10.0.4.4的路由。根据BGP选路原则，当BGP路由的下一跳不可达时，忽略此路由。

在R1上配置**next-hop-local**，再次在R3上查看该路由表。

```
[R1]bgp 64512
[R1-bgp]peer 10.0.3.3  next-hop-local
[R1-bgp]peer 10.0.2.2  next-hop-local
[R1-bgp]quit
[R3]display bgp routing-table

 BGP Local router ID is 10.0.3.3
 Status codes: * - valid, > - best, d - damped,
               h - history,  i - internal, s - suppressed, S - Stale
               Origin : i - IGP, e - EGP, ? - incomplete



 Total Number of Routes: 1
     Network          NextHop        MED        LocPrf    PrefVal Path/Ogn

 *>i  10.1.4.0/24       10.0.1.1       0          100        0      64513i
```

可以看到BGP路由10.1.4.0/24的下一跳为10.0.1.1，同时具有*号和>号，说明这条路由是正确并且最优的。

查看R3的路由表。

```
[R3]display ip routing-table
```

```
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Routing Tables: Public
        Destinations : 17      Routes : 17

Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface

      10.0.1.1/32   OSPF    10   3124        D   10.0.23.2       Serial2/0/0
      10.0.2.2/32   OSPF    10   1562        D   10.0.23.2       Serial2/0/0
      10.0.3.3/32   Direct  0    0           D   127.0.0.1       InLoopBack0
     10.0.12.0/24   OSPF    10   3124        D   10.0.23.2       Serial2/0/0
     10.0.23.0/24   Direct  0    0           D   10.0.23.3       Serial2/0/0
     10.0.23.2/32   Direct  0    0           D   10.0.23.2       Serial2/0/0
     10.0.23.3/32   Direct  0    0           D   127.0.0.1       InLoopBack0
   10.0.23.255/32   Direct  0    0           D   127.0.0.1       InLoopBack0
     10.0.35.0/24   Direct  0    0           D   10.0.35.3       Serial3/0/0
     10.0.35.3/32   Direct  0    0           D   127.0.0.1       InLoopBack0
     10.0.35.5/32   Direct  0    0           D   10.0.35.5       Serial3/0/0
   10.0.35.255/32   Direct  0    0           D   127.0.0.1       InLoopBack0
      10.1.4.0/24   IBGP    255  0           RD  10.0.1.1        Serial2/0/0
     127.0.0.0/8    Direct  0    0           D   127.0.0.1       InLoopBack0
     127.0.0.1/32   Direct  0    0           D   127.0.0.1       InLoopBack0
127.255.255.255/32  Direct  0    0           D   127.0.0.1       InLoopBack0
255.255.255.255/32  Direct  0    0           D   127.0.0.1       InLoopBack0
```

路由表出现路由10.1.4.0/24。

在R5上创建Loopback1，地址为10.1.5.5/24，发布进BGP，在R3上配置**next-hop-local**。

```
[R5]interface LoopBack 1
[R5-LoopBack1]ip address 10.1.5.5 24
[R5-LoopBack1]quit

[R5]bgp 64514
[R5-bgp]network 10.1.5.0 24

[R3]bgp 64512
[R3-bgp]peer 10.0.1.1 next-hop-local
[R3-bgp]peer 10.0.2.2 next-hop-local
```

在R4上查看是否学习到R5的Loopback 1连接网络的路由。分析**display bgp routing-table**的输出。

```
[R4]display bgp routing-table
```

```
BGP Local router ID is 10.0.4.4
Status codes: * - valid, > - best, d - damped,
              h - history,  i - internal, s - suppressed, S - Stale
              Origin : i - IGP, e - EGP, ? - incomplete


Total Number of Routes: 2
     Network            NextHop        MED        LocPrf    PrefVal Path/Ogn

*>   10.1.4.0/24        0.0.0.0         0                     0      i
*>   10.1.5.0/24        10.0.1.1                              0      64512 64514i
```

在R5上使用带源地址ping测试到R4的Loopback1地址的连通性。

```
[R5]ping -c 1 -a  10.1.5.5 10.1.4.4
  PING 10.1.4.4: 56  data bytes, press CTRL_C to break
    Reply from 10.1.4.4: bytes=56 Sequence=1 ttl=252 time=125 ms

  --- 10.1.4.4 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 125/125/125 ms
```

附加实验：思考并验证

什么情况下用物理地址直接建立EBGP邻居关系比较恰当？

为何默认发给EBGP邻居报文的TTL值为1？ 运行peer *group_name* ebgp-max-hop *[hop-count]*的默认值是多少？

最终设备配置

```
[R1]display current-configuration
[V200R001C00SPC200]
#
 sysname R1
#
 router id 10.0.1.1
```

```
#
interface Serial1/0/0
 link-protocol ppp
 ip address 10.0.12.1 255.255.255.0
#
interface Serial3/0/0
 link-protocol ppp
 ip address 10.0.14.1 255.255.255.0
#
interface LoopBack0
 ip address 10.0.1.1 255.255.255.255
#
bgp 64512
 timer keepalive 30 hold 90
 peer 10.0.2.2 as-number 64512
 peer 10.0.2.2 connect-interface LoopBack0
 peer 10.0.3.3 as-number 64512
 peer 10.0.3.3 connect-interface LoopBack0
 peer 10.0.4.4 as-number 64513
 peer 10.0.4.4 ebgp-max-hop 2
 peer 10.0.4.4 connect-interface LoopBack0
 #
 ipv4-family unicast
  undo synchronization
  peer 10.0.2.2 enable
  peer 10.0.2.2 next-hop-local
  peer 10.0.3.3 enable
  peer 10.0.3.3 next-hop-local
  peer 10.0.4.4 enable
#
ospf 1
 area 0.0.0.0
  network 10.0.12.0 0.0.0.255
  network 10.0.1.1 0.0.0.0
#
 ip route-static 10.0.4.4 255.255.255.255 10.0.14.4
return
```

```
[R2]display current-configuration
[V200R001C00SPC200]
#
 sysname R2
#
```

```
 router id 10.0.2.2
#
interface Serial1/0/0
 link-protocol ppp
 ip address 10.0.12.2 255.255.255.0
#
interface Serial2/0/0
 link-protocol ppp
 ip address 10.0.23.2 255.255.255.0
#
interface LoopBack0
 ip address 10.0.2.2 255.255.255.255
#
bgp 64512
 peer 10.0.1.1 as-number 64512
 peer 10.0.1.1 connect-interface LoopBack0
 peer 10.0.3.3 as-number 64512
 peer 10.0.3.3 connect-interface LoopBack0
 #
 ipv4-family unicast
  undo synchronization
  peer 10.0.1.1 enable
  peer 10.0.3.3 enable
#
ospf 1
 area 0.0.0.0
  network 10.0.12.0 0.0.0.255
  network 10.0.23.0 0.0.0.255
  network 10.0.2.2 0.0.0.0
return
```

[R3]display current-configuration

```
[V200R001C00SPC200]
#
 sysname R3
#
 router id 10.0.3.3
#
interface Serial2/0/0
 link-protocol ppp
 ip address 10.0.23.3 255.255.255.0
#
interface Serial3/0/0
```

```
 link-protocol ppp
 ip address 10.0.35.3 255.255.255.0
#
interface LoopBack0
 ip address 10.0.3.3 255.255.255.255
#
bgp 64512
 peer 10.0.1.1 as-number 64512
 peer 10.0.1.1 connect-interface LoopBack0
 peer 10.0.2.2 as-number 64512
 peer 10.0.2.2 connect-interface LoopBack0
 peer 10.0.35.5 as-number 64514
 #
 ipv4-family unicast
  undo synchronization
  peer 10.0.1.1 enable
  peer 10.0.1.1 next-hop-local
  peer 10.0.2.2 enable
  peer 10.0.2.2 next-hop-local
  peer 10.0.35.5 enable
#
ospf 1
 area 0.0.0.0
  network 10.0.23.0 0.0.0.255
  network 10.0.3.3 0.0.0.0
return
```

```
[R4]display current-configuration
[V200R001C00SPC200]
#
 sysname R4
#
interface Serial1/0/0
 link-protocol ppp
 ip address 10.0.14.4 255.255.255.0
#
interface LoopBack0
 ip address 10.0.4.4 255.255.255.255
#
interface LoopBack1
 ip address 10.1.4.4 255.255.255.0
#
bgp 64513
```

```
 peer 10.0.1.1 as-number 64512
 peer 10.0.1.1 ebgp-max-hop 2
 peer 10.0.1.1 connect-interface LoopBack0
 #
 ipv4-family unicast
  undo synchronization
  network 10.0.4.0 255.255.255.0
  network 10.1.4.0 255.255.255.0
  peer 10.0.1.1 enable
#
 ip route-static 10.0.1.1 255.255.255.255 10.0.14.1
return
```

```
[R5]display current-configuration
[V200R001C00SPC200]
#
 sysname R5
#
interface Serial1/0/0
 link-protocol ppp
 ip address 10.0.35.5 255.255.255.0
#
interface LoopBack0
 ip address 10.0.5.5 255.255.255.255
#
interface LoopBack1
 ip address 10.1.5.5 255.255.255.0
#
bgp 64514
 peer 10.0.35.3 as-number 64512
 #
 ipv4-family unicast
  undo synchronization
  network 10.1.5.0 255.255.255.0
  peer 10.0.35.3 enable
return
```

实验 3-2 BGP 路由汇总

学习目的

- 掌握使用**network**命令实现的BGP路由汇总的配置方法
- 掌握实现抑制具体路由的BGP路由汇总的配置方法
- 掌握改变汇总路由属性的配置方法
- 掌握在路由汇总时使用AS-SET的配置方法

拓扑图

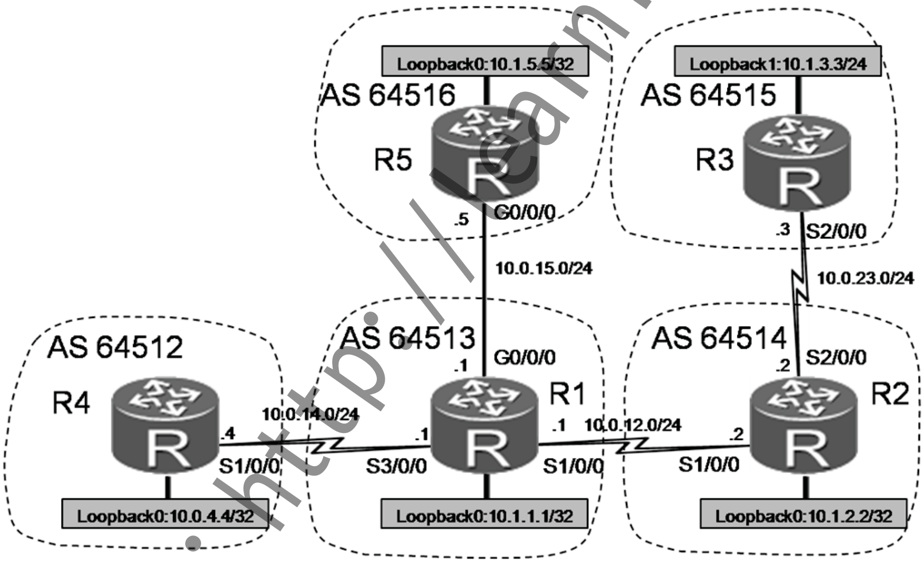

图3-2 BGP路由汇总

场景

你是公司的网络管理员。公司的网络采用了BGP协议作为路由协议。公司的网络由多个自制系统组成，不同的分支机构使用了不同的AS号。随着公司规模的扩大，路由器中已经有越来越多的路由表，进行BGP的路由汇总迫在眉睫。你测试了几种进行路由汇总的方法，最终选择了合适的方式实现了路由汇总。

学习任务

步骤一. 基础配置与 IP 编址

给所有路由器配置物理接口及Loopback接口的IP地址和掩码。注意各Loopback接口地址均为32位掩码。

```
<R1>system-view
Enter system view, return user view with Ctrl+Z.
[R1]interface Serial 1/0/0
[R1-Serial1/0/0]ip address 10.0.12.1 255.255.255.0
[R1-Serial1/0/0]interface Serial 3/0/0
[R1-Serial3/0/0]ip address 10.0.14.1 255.255.255.0
[R1-Serial3/0/0]interface GigabitEthernet 0/0/0
[R1-GigabitEthernet0/0/0]ip add 10.0.15.1 255.255.255.0
[R1-GigabitEthernet0/0/0]interface  LoopBack 0
[R1-LoopBack0]ip address 10.1.1.1 255.255.255.255

<R2>system-view
Enter system view, return user view with Ctrl+Z.
[R2]interface Serial 1/0/0
[R2-Serial1/0/0]ip address 10.0.12.2 255.255.255.0
[R2-Serial1/0/0]interface Serial 2/0/0
[R2-Serial2/0/0]ip address 10.0.23.2 255.255.255.0
[R2-Serial2/0/0]interface loopback 0
[R2-LoopBack0]ip address 10.1.2.2 255.255.255.255

<R3>system-view
Enter system view, return user view with Ctrl+Z.
[R3]interface Serial 2/0/0
[R3-Serial2/0/0]ip address 10.0.23.3 255.255.255.0
[R3-Serial2/0/0]interface loopback 1
[R3-LoopBack1]ip address 10.1.3.3 255.255.255.255

<R4>system-view
Enter system view, return user view with Ctrl+Z.
[R4]interface Serial 1/0/0
[R4-Serial1/0/0]ip address 10.0.14.4 255.255.255.0
[R4-Serial1/0/0]interface loopback 0
[R4-LoopBack0]ip address 10.0.4.4 255.255.255.255

<R5>system-view
```

```
Enter system view, return user view with Ctrl+Z.
[R5]interface GigabitEthernet 0/0/0
[R5-GigabitEthernet0/0/0]ip address 10.0.15.5 255.255.255.0
[R5-GigabitEthernet0/0/0]interface loopback 0
[R5-LoopBack0]ip address 10.1.5.5 255.255.255.255
```

配置完成后，测试直连链路的连通性。

```
[R1]ping -c 1 10.0.12.2
  PING 10.0.12.2: 56  data bytes, press CTRL_C to break
    Reply from 10.0.14.4: bytes=56 Sequence=1 ttl=255 time=41 ms

  --- 10.0.12.2 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 41/41/41 ms

[R1]ping -c 1 10.0.14.4
  PING 10.0.14.4: 56  data bytes, press CTRL_C to break
    Reply from 10.0.14.4: bytes=56 Sequence=1 ttl=255 time=41 ms

  --- 10.0.14.4 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 41/41/41 ms

[R1]ping -c 1 10.0.15.5
  PING 10.0.15.5: 56  data bytes, press CTRL_C to break
    Reply from 10.0.12.2: bytes=56 Sequence=1 ttl=255 time=34 ms

  --- 10.0.15.5 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 34/34/34 ms

[R1]ping -c 1 10.0.23.3
  PING 10.0.23.3: 56  data bytes, press CTRL_C to break
    Reply from 10.0.12.2: bytes=56 Sequence=1 ttl=255 time=34 ms

  --- 10.0.23.3 ping statistics ---
    1 packet(s) transmitted
```

```
1 packet(s) received
0.00% packet loss
round-trip min/avg/max = 34/34/34 ms
```

步骤二.　配置 EBGP 及发布路由

各直连路由器之间直接使用物理接口地址建立BGP对等体关系。

```
[R1]bgp 64513
[R1-bgp]peer 10.0.12.2 as-number 64514
[R1-bgp]peer 10.0.14.4 as-number 64512
[R1-bgp]peer 10.0.15.5 as-number 64516

[R2]bgp 64514
[R2-bgp]peer 10.0.12.1 as-number 64513
[R2-bgp]peer 10.0.23.3 as-number 64515

[R3]bgp 64515
[R3-bgp]peer 10.0.23.2 as-number 64514

[R4]bgp 64512
[R4-bgp]peer 10.0.14.1 as-number 64513

[R5]bgp 64516
[R5-bgp]peer 10.0.15.1 as-number 64513
```

配置完成以后检查对等体关系，

```
[R1]display bgp peer
 BGP local router ID : 10.0.12.1
 Local AS number : 64513
 Total number of peers : 3              Peers in established state : 3

 Peer            V    AS  MsgRcvd  MsgSent  OutQ   Up/Down     State      PrefRcv

  10.0.12.2       4    64514       3        3     0 00:01:55 Established        0
  10.0.14.4       4    64512       3        3     0 00:01:03 Established        0
  10.0.15.5       4    64516       2        3     0 00:00:03 Established        0

[R2]display bgp peer

 BGP local router ID : 10.0.12.2
 Local AS number : 64514
```

```
 Total number of peers : 2              Peers in established state : 2

  Peer             V        AS  MsgRcvd  MsgSent  OutQ  Up/Down       State PrefRcv

  10.0.12.1        4      64513        4        6     0 00:02:51 Established       0
  10.0.23.3        4      64515        6        6     0 00:04:42 Established       0
[R3]display bgp peer

 BGP local router ID : 10.0.23.3
 Local AS number : 64515
 Total number of peers : 1              Peers in established state : 1

  Peer             V        AS  MsgRcvd  MsgSent  OutQ  Up/Down       State PrefRcv

  10.0.23.2        4      64514        6        6     0 00:04:54 Established       0
[R4]display bgp peer

 BGP local router ID : 10.0.14.4
 Local AS number : 64512
 Total number of peers : 1              Peers in established state : 1

  Peer             V        AS  MsgRcvd  MsgSent  OutQ  Up/Down      State   PrefRcv

  10.0.14.1        4      64513        4        4     0 00:02:46 Established       0
[R5]display bgp peer

 BGP local router ID : 10.0.15.5
 Local AS number : 64516
 Total number of peers : 1              Peers in established state : 1

  Peer             V        AS  MsgRcvd  MsgSent  OutQ  Up/Down      State   PrefRcv

  10.0.15.1        4      64513        3        3     0 00:01:26 Established       0
```

BGP邻居状态此时全部都是Established状态。

使用**network**命令将各个路由器的Loopback接口网段发布进BGP。

```
[R1]bgp 64513
[R1-bgp]network 10.1.1.1 255.255.255.255
```

```
[R2]bgp 64514
[R2-bgp]network 10.1.2.2 255.255.255.255

[R3]bgp 64515
[R3-bgp]network 10.1.3.3 255.255.255.255

[R4]bgp 64512
[R4-bgp]network 10.0.4.4 255.255.255.255

[R5]bgp 64516
[R5-bgp]network 10.1.5.5 255.255.255.255
```

在R4上查看BGP路由表，观察AS-PATH属性的信息。

```
[R4]display bgp routing-table

 BGP Local router ID is 10.0.14.4
 Status codes: * - valid, > - best, d - damped,
               h - history,  i - internal, s - suppressed, S - Stale
               Origin : i - IGP, e - EGP, ? - incomplete


 Total Number of Routes: 5
      Network            NextHop        MED        LocPrf    PrefVal Path/Ogn

 *>   10.0.4.4/32        0.0.0.0         0                     0      i
 *>   10.1.1.1/32        10.0.14.1       0                     0      64513i
 *>   10.1.2.2/32        10.0.14.1                             0      64513 64514i
 *>   10.1.3.3/32        10.0.14.1                             0      64513 64514
64515i
 *>   10.1.5.5/32        10.0.14.1                             0      64513 64516i
```

步骤三.　使用 Network 实现对 BGP 路由的汇总

现在需要在R1上进行路由汇总。

在R1上添加指向Null0接口的静态路由10.1.0.0/16，并使用**network**命令发布该路由。

```
[R1]ip route-static 10.1.0.0 16 NULL 0
[R1]bgp 64513
[R1-bgp]network 10.1.0.0 255.255.0.0
```

在R4上查看路由表，观察汇总路由是否存在。

```
<R4>display bgp routing-table

 BGP Local router ID is 10.0.14.4
 Status codes: * - valid, > - best, d - damped,
               h - history,  i - internal, s - suppressed, S - Stale
               Origin : i - IGP, e - EGP, ? - incomplete


 Total Number of Routes: 6
      Network            NextHop        MED        LocPrf    PrefVal Path/Ogn

 *>   10.0.4.4/32         0.0.0.0         0                     0       i
 *>   10.1.0.0/16         10.0.14.1       0                     0       64513i
 *>   10.1.1.1/32         10.0.14.1       0                     0       64513i
 *>   10.1.2.2/32         10.0.14.1                             0      64513 64514i
 *>   10.1.3.3/32         10.0.14.1                             0      64513 64514
64515i
 *>   10.1.5.5/32         10.0.14.1                             0      64513 64516i
```

设置名为**pref_detail_control**的前缀列表，对向对等体R4发送的路由进行过滤，不允许汇总路由中包括的详细路由被发送过去。

```
[R1]ip ip-prefix pref_detail_control index 10 permit 10.1.0.0 8 less-equal 24
[R1]bgp 64513
[R1-bgp]peer 10.0.14.4 ip-prefix pref_detail_control export
```

在R4上再次查看BGP路由表。注意观察汇总后路由的As-path属性。

```
<R4>display bgp routing-table

 BGP Local router ID is 10.0.14.4
 Status codes: * - valid, > - best, d - damped,
               h - history,  i - internal, s - suppressed, S - Stale
               Origin : i - IGP, e - EGP, ? - incomplete


 Total Number of Routes: 2
      Network            NextHop        MED        LocPrf    PrefVal Path/Ogn

 *>   10.0.4.4/32         0.0.0.0         0                     0       i
 *>   10.1.0.0/16         10.0.14.1       0                     0       64513i
```

步骤四.　使用 Aggregate 实现对 BGP 路由的汇总

删去步骤三中使用的前缀列表及**network**命令发布的汇总路由。

使用**Aggregate**命令对路由10.1.0.0/16进行汇总，使用默认方式进行。

```
[R1]undo ip ip-prefix pref_detail_control
[R1]bgp 64513
[R1-bgp]undo network 10.1.0.0 255.255.0.0
[R1-bgp]undo peer 10.0.14.4 ip-prefix pref_detail_control export
[R1-bgp]aggregate 10.1.0.0 255.255.0.0
```

在R1、R4查看BGP路由表。观察汇总路由的Origin属性。

```
[R1]display bgp routing-table

 BGP Local router ID is 10.0.12.1
 Status codes: * - valid, > - best, d - damped,
               h - history,  i - internal, s - suppressed, S - Stale
               Origin : i - IGP, e - EGP, ? - incomplete


 Total Number of Routes: 6
     Network            NextHop        MED        LocPrf    PrefVal Path/Ogn

 *>   10.0.4.4/32        10.0.14.4      0                     0      64512i
 *>   10.1.0.0/16        127.0.0.1                            0      i
 *>   10.1.1.1/32        0.0.0.0        0                     0      i
 *>   10.1.2.2/32        10.0.12.2      0                     0      64514i
 *>   10.1.3.3/32        10.0.12.2                            0      64514 64515i
 *>   10.1.5.5/32        10.0.15.5      0                     0      64516i

<R4>display bgp routing-table

 BGP Local router ID is 10.0.14.4
 Status codes: * - valid, > - best, d - damped,
               h - history,  i - internal, s - suppressed, S - Stale
               Origin : i - IGP, e - EGP, ? - incomplete


 Total Number of Routes: 6
     Network            NextHop        MED        LocPrf    PrefVal Path/Ogn

 *>   10.0.4.4/32        0.0.0.0        0                     0      i
```

```
*>   10.1.0.0/16        10.0.14.1                            0      64513i
*>   10.1.1.1/32        10.0.14.1       0                    0      64513i
*>   10.1.2.2/32        10.0.14.1                            0      64513 64514i
*>   10.1.3.3/32        10.0.14.1                            0      64513 64514
64515i
*>   10.1.5.5/32        10.0.14.1                            0      64513 64516i
```

汇总路由的**origin**属性没有改变，仍旧是IGP。

在R1配置路由汇总时，抑制明细路由，仅通告聚合路由。

```
[R1-bgp]aggregate 10.1.0.0 255.255.0.0 detail-suppressed
```

在R4查看BGP路由表。

```
[R4]display bgp routing-table

 BGP Local router ID is 10.0.14.4
 Status codes: * - valid, > - best, d - damped,
               h - history,  i - internal, s - suppressed, S - Stale
               Origin : i - IGP, e - EGP, ? - incomplete


 Total Number of Routes: 2
      Network            NextHop        MED        LocPrf    PrefVal Path/Ogn

 *>   10.0.4.4/32        0.0.0.0         0                     0      i
 *>   10.1.0.0/16        10.0.14.1                             0      64513i
```

这时R4上已经看不到明细路由。

观察R1的全局路由表，查看路由10.1.0.0/16的下一跳。

```
[R1]display ip routing-table
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Routing Tables: Public
         Destinations : 21      Routes : 21

Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface

      10.0.4.4/32   EBGP    255  0           D   10.0.14.4       Serial3/0/0
     10.0.12.0/24   Direct  0    0           D   10.0.12.1       Serial1/0/0
     10.0.12.1/32   Direct  0    0           D   127.0.0.1       InLoopBack0
     10.0.12.2/32   Direct  0    0           D   10.0.12.2       Serial1/0/0
   10.0.12.255/32   Direct  0    0           D   127.0.0.1       InLoopBack0
```

```
     10.0.14.0/24   Direct  0    0           D  10.0.14.1      Serial3/0/0
     10.0.14.1/32   Direct  0    0           D  127.0.0.1      InLoopBack0
     10.0.14.4/32   Direct  0    0           D  10.0.14.4      Serial3/0/0
   10.0.14.255/32   Direct  0    0           D  127.0.0.1      InLoopBack0
     10.0.15.0/24   Direct  0    0           D  10.0.15.1   GigabitEthernet0/0/0
     10.0.15.1/32   Direct  0    0           D  127.0.0.1      InLoopBack0
   10.0.15.255/32   Direct  0    0           D  127.0.0.1      InLoopBack0
      10.1.0.0/16   Static  60   0           D  0.0.0.0        NULL0
      10.1.1.1/32   Direct  0    0           D  127.0.0.1      InLoopBack0
      10.1.2.2/32   EBGP    255  0           D  10.0.12.2      Serial1/0/0
      10.1.3.3/32   EBGP    255  0           D  10.0.12.2      Serial1/0/0
      10.1.5.5/32   EBGP    255  0           D  10.0.15.5   GigabitEthernet0/0/0
     127.0.0.0/8    Direct  0    0           D  127.0.0.1      InLoopBack0
     127.0.0.1/32   Direct  0    0           D  127.0.0.1      InLoopBack0
127.255.255.255/32  Direct  0    0           D  127.0.0.1      InLoopBack0
255.255.255.255/32  Direct  0    0           D  127.0.0.1      InLoopBack0
```

因为这条聚合路由是R1上配置的，所以出接口为Null0，这样定义的汇总路由有利于避免路由环路的产生。

观察R1的BGP路由表，观察明细路由。

```
[R1]display bgp routing-table

 BGP Local router ID is 10.0.12.1
 Status codes: * - valid, > - best, d - damped,
               h - history,  i - internal, s - suppressed, S - Stale
               Origin : i - IGP, e - EGP, ? - incomplete


 Total Number of Routes: 6
      Network            NextHop        MED        LocPrf    PrefVal Path/Ogn

 *>   10.0.4.4/32        10.0.14.4       0                     0      64512i
 *>   10.1.0.0/16        127.0.0.1                             0      i
 s>   10.1.1.1/32        0.0.0.0         0                     0      i
 s>   10.1.2.2/32        10.0.12.2       0                     0      64514i
 s>   10.1.3.3/32        10.0.12.2                             0      64514 64515i
 s>   10.1.5.5/32        10.0.15.5       0                     0      64516i
```

使用了**detail-suppressed**参数，对外仅发送汇总路由，这时在明细路由条目前面多了个标识s，代表在路由汇总时，明细路由被抑制。

步骤五. 改变汇总路由的属性

缺省情况下，BGP不将团体属性发布给任何对等体。

配置R5向R1、R1向R4通告团体属性。

```
[R5]bgp 64516
[R5-bgp]peer 10.0.15.1 advertise-community

[R1]bgp 64513
[R1-bgp]peer 10.0.14.4 advertise-community
```

验证进行路由汇总后团体属性会丢失。

在R5上对R5通告的10.1.5.5/32加上100的团体属性，并向R1通告。

```
[R5]acl number 2000
[R5-acl-basic-2000]rule 0 permit source 10.1.5.5 0
[R5-acl-basic-2000]route-policy set_comm permit node 10
[R5-route-policy]if-match acl 2000
[R5-route-policy]apply community 100
[R5-route-policy]bgp 64516
[R5-bgp]peer 10.0.15.1 route-policy set_comm export
```

在R1上查看该路由是否携带该团体属性。

```
<R1>display bgp routing-table community

 BGP Local router ID is 10.0.14.4
 Status codes: * - valid, > - best, d - damped,
               h - history,  i - internal, s - suppressed, S - Stale
               Origin : i - IGP, e - EGP, ? - incomplete



 Total Number of Routes: 2
     Network            NextHop        MED        LocPrf    PrefVal Community

 *>   10.1.5.0/24        10.0.15.5                              0      <0:100>
```

在R4上查看汇总后路由是否携带该属性。

```
<R4>display bgp routing-table community

 Total Number of Routes: 0
```

R4没有任何携带团体属性的路由。

在R1上配置路由策略**add_comm**，将100:2的团体属性加到汇总后的路由上。

```
[R1]acl number 2000
[R1-acl-basic-2000]rule 0 permit source 10.1.0.0 0.0.255.255
[R1-acl-basic-2000]route-policy add_comm permit node 10
[R1-route-policy]if-match acl 2000
[R1-route-policy]apply community 100:2
[R1-route-policy]bgp 64513
[R1-bgp]aggregate 10.1.0.0 255.255.0.0 attribute-policy add_comm
```

在R4上观察该汇总路由是否携带100:2团体属性。

```
<R4>display bgp routing-table community

 BGP Local router ID is 10.0.14.4
 Status codes: * - valid, > - best, d - damped,
               h - history,  i - internal, s - suppressed, S - Stale
               Origin : i - IGP, e - EGP, ? - incomplete


 Total Number of Routes: 1
      Network            NextHop        MED        LocPrf    PrefVal Community

 *>   10.1.0.0/16        10.0.14.1                            0      <100:2>
```

R4学习到的汇总路由携带该属性。

步骤六.　使用 AS-SET 配置 AS-PATH 属性

路由在汇总后，默认会丢弃AS-PATH属性信息，AS-PATH的丢失可能会造成环路。为避免信息丢失带来的风险，汇总路由可以在汇总时添加AS-Set属性。

配置R1在执行路由汇总时添加AS-Set属性。

```
 [R1]bgp 64513
[R1-bgp]aggregate 10.1.0.0 255.255.0.0 detail-suppressed as-set
```

观察R1、R4中BGP路由表中汇总路由的AS-PATH属性信息。

```
[R1]display bgp routing-table

 BGP Local router ID is 10.0.12.1
```

```
 Status codes: * - valid, > - best, d - damped,
               h - history,  i - internal, s - suppressed, S - Stale
               Origin : i - IGP, e - EGP, ? - incomplete

 Total Number of Routes: 6
      Network            NextHop        MED        LocPrf    PrefVal Path/Ogn

 *>   10.0.4.4/32        10.0.14.4       0                     0      64512i
 *>   10.1.0.0/16        127.0.0.1                             0      {64514 64515
64516}i
 s>   10.1.1.1/32        0.0.0.0         0                     0      i
 s>   10.1.2.2/32        10.0.12.2       0                     0      64514i
 s>   10.1.3.3/32        10.0.12.2                             0      64514 64515i
 s>   10.1.5.5/32        10.0.15.5       0                     0      64516i

<R4>display bgp routing-table

 BGP Local router ID is 10.0.14.4
 Status codes: * - valid, > - best, d - damped,
               h - history,  i - internal, s - suppressed, S - Stale
               Origin : i - IGP, e - EGP, ? - incomplete


 Total Number of Routes: 2
      Network            NextHop        MED        LocPrf    PrefVal Path/Ogn

 *>   10.0.4.4/32        0.0.0.0         0                     0      i
 *>   10.1.0.0/16        10.0.14.1                             0      64513 {64514
64515 64516}i
```

加了AS-Set属性之后的汇总路由的AS-PATH包含了具体路由的AS路径信息。

在R3上停止通告10.1.3.3/32，重置对等体关系。

```
[R3]bgp 64515
[R3-bgp]undo network 10.1.3.3 255.255.255.255
[R3-bgp]return
<R3>reset bgp all
```

带邻居关系重新建立后，观察R4上学习到的汇总路由的AS-PATH属性信息。

```
<R4>display bgp routing-table
```

```
 BGP Local router ID is 10.0.14.4
 Status codes: * - valid, > - best, d - damped,
            h - history,  i - internal, s - suppressed, S - Stale
            Origin : i - IGP, e - EGP, ? - incomplete


 Total Number of Routes: 2
     Network          NextHop       MED       LocPrf    PrefVal Path/Ogn

 *>   10.0.4.4/32       0.0.0.0       0                   0      i
 *>   10.1.0.0/16       10.0.14.1                         0      64513 {64514
64516}i
```

这时发现AS-PATH属性中已经没有AS号码64515。

附加实验：思考并验证

这个例子中，完成步骤六以后，R5是否能访问到R3的Loopback地址？

Aggregate和**Summary automatic**有什么区别？

最终设备配置

```
<R1>display current-configuration
[V200R001C00SPC200]
#
 sysname R1
#
acl number 2000
 rule 0 permit source 10.1.0.0 0.0.255.255
#
interface Serial1/0/0
 link-protocol ppp
 ip address 10.0.12.1 255.255.255.0
#
interface Serial3/0/0
 link-protocol ppp
 ip address 10.0.14.1 255.255.255.0
#
interface GigabitEthernet0/0/0
 ip address 10.0.15.1 255.255.255.0
```

```
#
interface LoopBack0
 ip address 10.1.1.1 255.255.255.255
#
bgp 64513
 peer 10.0.12.2 as-number 64514
 peer 10.0.14.4 as-number 64512
 peer 10.0.15.5 as-number 64516
 #
 ipv4-family unicast
  undo synchronization
  aggregate 10.1.0.0 255.255.0.0 as-set detail-suppressed
  network 10.1.1.1 255.255.255.255
  peer 10.0.12.2 enable
  peer 10.0.14.4 enable
  peer 10.0.14.4 advertise-community
  peer 10.0.15.5 enable
#
route-policy add_comm permit node 10
 if-match acl 2000
 apply community 100:2
#
 ip route-static 10.1.0.0 255.255.0.0 NULL0
#
return

<R2>display current-configuration
[V200R001C00SPC200]
#
 sysname R2
#
interface Serial1/0/0
 link-protocol ppp
 ip address 10.0.12.2 255.255.255.0
#
interface Serial2/0/0
 link-protocol ppp
 ip address 10.0.23.2 255.255.255.0
#
interface LoopBack0
 ip address 10.1.2.2 255.255.255.255
#
bgp 64514
```

```
 peer 10.0.12.1 as-number 64513
 peer 10.0.23.3 as-number 64515
 #
 ipv4-family unicast
  undo synchronization
  network 10.1.2.2 255.255.255.255
  peer 10.0.12.1 enable
  peer 10.0.23.3 enable
#
return

<R3>display current-configuration
[V200R001C00SPC200]
#
 sysname R3
#
interface Serial2/0/0
 link-protocol ppp
 ip address 10.0.23.3 255.255.255.0
#
interface LoopBack1
 ip address 10.1.3.3 255.255.255.255
#
bgp 64515
 peer 10.0.23.2 as-number 64514
 #
 ipv4-family unicast
  undo synchronization
  peer 10.0.23.2 enable
#
return

<R4>display current-configuration
[V200R001C00SPC200]
#
 sysname R4
#
interface Serial1/0/0
 link-protocol ppp
 ip address 10.0.14.4 255.255.255.0
#
interface LoopBack0
 ip address 10.0.4.4 255.255.255.255
```

```
#
bgp 64512
 peer 10.0.14.1 as-number 64513
 #
 ipv4-family unicast
  undo synchronization
  network 10.0.4.4 255.255.255.255
  peer 10.0.14.1 enable
#
return
```

```
<R5>display current-configuration
[V200R001C00SPC200]
#
 sysname R5
#
acl number 2000
 rule 0 permit source 10.1.5.5 0
#
interface GigabitEthernet0/0/0
 ip address 10.0.15.5 255.255.255.0
#
interface LoopBack0
 ip address 10.1.5.5 255.255.255.255
#
bgp 64516
 peer 10.0.15.1 as-number 64513
 #
 ipv4-family unicast
  undo synchronization
  network 10.1.5.5 255.255.255.255
  peer 10.0.15.1 enable
  peer 10.0.15.1 advertise-community
  peer 10.0.15.1 route-policy set_comm export
#
route-policy set_comm permit node 10
 if-match acl 2000
 apply community 100
#
return
```

实验 3-3 BGP 属性与路径选择 1

学习目的

- 掌握通过配置AS-Path属性来影响路径选择的方法
- 掌握通过修改Origin属性来影响路径选择的配置方法
- 掌握通过修改Local-Pref属性来影响路径选择的配置方法
- 掌握通过修改MED属性来影响路径选择的配置方法

拓扑图

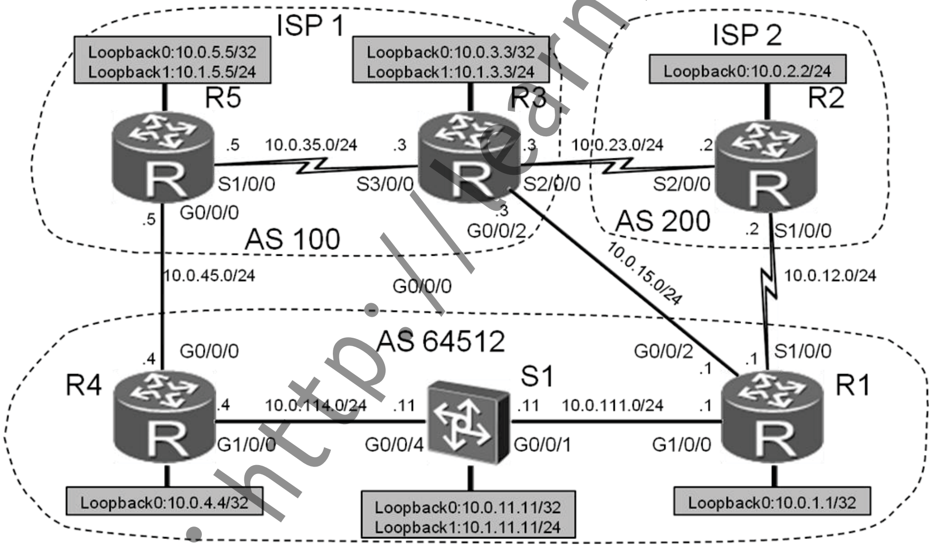

图3-3 BGP属性与路径选择

场景

你是公司的网络管理员。公司的网络采用了BGP协议接入了两个服务运营商。公司自己采用了私有的AS号64512，ISP1的AS号为100，公司共有2条链路接入ISP1。ISP2的AS号为200，公司租用了一条线路接入ISP2。现在Internet上的部分用户反应访问公司网站的速度较慢，你通过改变BGP的各种属性达到了调整路

由走向的目的。

学习任务

步骤一. 基础配置与 IP 编址

给所有路由器配置物理接口及Loopback接口的IP地址和掩码。注意各Loopback 0接口均使用32位掩码。

```
<Huawei>system-view
Enter system view, return user view with Ctrl+Z.
[Huawei]sysname R1
[R1]interface Serial 1/0/0
[R1-Serial1/0/0]ip address 10.0.12.1 255.255.255.0
[R1-Serial1/0/0]interface GigabitEthernet 0/0/2
[R1-GigabitEthernet0/0/2]ip address 10.0.15.1 255.255.255.0
[R1-GigabitEthernet0/0/2]interface GigabitEthernet 0/0/1
[R1-GigabitEthernet0/0/1]ip address 10.0.111.1 255.255.255.0
[R1-GigabitEthernet0/0/1]interface LoopBack 0
[R1-LoopBack0]ip address 10.0.1.1 255.255.255.255

<Huawei>system-view
Enter system view, return user view with Ctrl+Z.
[Huawei]sysname R2
[R2]interface Serial 1/0/0
[R2-Serial1/0/0]ip address 10.0.12.2 255.255.255.0
[R2-Serial1/0/0]int Serial 2/0/0
[R2-Serial2/0/0]ip address 10.0.23.2 255.255.255.0
[R2-Serial2/0/0]interface LoopBack 0
[R2-LoopBack0]ip address 10.0.2.2 255.255.255.0

<Huawei>system-view
Enter system view, return user view with Ctrl+Z.
[Huawei]sysname R3
[R3]interface GigabitEthernet 0/0/2
[R3-GigabitEthernet0/0/2]ip address 10.0.15.3 255.255.255.0
[R3-GigabitEthernet0/0/2]interface Serial 2/0/0
[R3-Serial2/0/0]ip address 10.0.23.3 255.255.255.0
[R3-Serial2/0/0]interface Serial 3/0/0
[R3-Serial3/0/0]ip address 10.0.35.3 255.255.255.0
[R3-Serial3/0/0]interface loopback 0
[R3-LoopBack0]ip address 10.0.3.3 255.255.255.255
```

```
<Huawei>system-view
Enter system view, return user view with Ctrl+Z.
[Huawei]sysname R4
[R4]interface GigabitEthernet 0/0/1
[R4-GigabitEthernet0/0/1]ip address 10.0.114.4 255.255.255.0
[R4-GigabitEthernet0/0/1]interface GigabitEthernet 0/0/0
[R4-GigabitEthernet0/0/0]ip address 10.0.45.4 255.255.255.0
[R4-GigabitEthernet0/0/0]interface loopback 0
[R4-LoopBack0]ip address 10.0.4.4 255.255.255.255

<Huawei>system-view
Enter system view, return user view with Ctrl+Z.
[Huawei]sysname R5
[R5]interface Serial 1/0/0
[R5-Serial1/0/0]ip address 10.0.35.5 255.255.255.0
[R5-Serial1/0/0]interface GigabitEthernet 0/0/0
[R5-GigabitEthernet0/0/0]ip address 10.0.45.5 255.255.255.0
[R5-GigabitEthernet0/0/0]interface loopback 0
[R5-LoopBack0]ip address 10.0.5.5 255.255.255.255
```

配置完成后，测试直连链路的连通性。

```
<R1>ping -c 1 10.0.12.2
  PING 10.0.12.2: 56  data bytes, press CTRL_C to break
    Reply from 10.0.12.2: bytes=56 Sequence=1 ttl=255 time=29 ms

  --- 10.0.12.2 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 29/29/29 ms

[R1]ping -c 1 10.0.15.3
  PING 10.0.15.3: 56  data bytes, press CTRL_C to break
    Reply from 10.0.15.3: bytes=56 Sequence=1 ttl=255 time=59 ms

  --- 10.0.15.3 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 59/59/59 ms

<R2>ping -c 1 10.0.23.3
```

```
  PING 10.0.23.3: 56  data bytes, press CTRL_C to break
    Reply from 10.0.23.3: bytes=56 Sequence=1 ttl=255 time=32 ms

  --- 10.0.23.3 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 32/32/32 ms

[R3]ping -c 1 10.0.35.5
  PING 10.0.35.5: 56  data bytes, press CTRL_C to break
    Reply from 10.0.35.5: bytes=56 Sequence=1 ttl=255 time=36 ms

  --- 10.0.35.5 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 36/36/36 ms

<R4>ping -c 1 10.0.45.5
  PING 10.0.45.5: 56  data bytes, press CTRL_C to break
    Reply from 10.0.45.5: bytes=56 Sequence=1 ttl=255 time=11 ms

  --- 10.0.45.5 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 11/11/11 ms
```

步骤二.　配置 IGP 及 BGP

在AS 64512内部使用OSPF作为IGP，所有设备属于区域0。

R1的G0/0/1和Loopback 0连接的网段运行OSPF。

```
[R1]ospf
[R1-ospf-1]area 0
[R1-ospf-1-area-0.0.0.0]network 10.0.111.1 0.0.0.0
[R1-ospf-1-area-0.0.0.0]network 10.0.1.1 0.0.0.0
```

在S1上创建Vlan 111，配置Vlanif地址与R1进行互联。

创建Vlan114，配置Vlanif地址与R4进行互联。

互联接口使用Access模式，S1的Vlanif 111、Vlanif 114 和Loopback 0连接的网段运行OSPF。

```
[S1]vlan 111
[S1-vlan111]vlan 114
[S1]interface vlan 111
[S1-Vlanif111]ip address 10.0.111.11 255.255.255.0
[S1-Vlanif111]int vlan 114
[S1-Vlanif114]ip address 10.0.114.11 255.255.255.0
[S1]interface loopback 0
[S1-LoopBack0]ip address 10.0.11.11 255.255.255.255
[S1-LoopBack0]interface GigabitEthernet 0/0/1
[S1-GigabitEthernet0/0/1]port link-type access
[S1-GigabitEthernet0/0/1]port default vlan 111
[S1-GigabitEthernet0/0/1]interface GigabitEthernet 0/0/4
[S1-GigabitEthernet0/0/4]port link-type access
[S1-GigabitEthernet0/0/4]port default vlan 114
[S1-GigabitEthernet0/0/4]ospf
[S1-ospf-1]area 0
[S1-ospf-1-area-0.0.0.0]network 10.0.111.11 0.0.0.0
[S1-ospf-1-area-0.0.0.0]network 10.0.114.11 0.0.0.0
[S1-ospf-1-area-0.0.0.0]network 10.0.11.11 0.0.0.0
```

R4的G0/0/1和Loopback 0连接的网段运行OSPF，

```
[R4]ospf
[R4-ospf-1]area 0
[R4-ospf-1-area-0.0.0.0]network 10.0.114.4 0.0.0.0
[R4-ospf-1-area-0.0.0.0]network 10.0.4.4 0.0.0.0
```

检查是否学习到其他设备的Loopback 0地址所在的网段。

```
[R1]display ip routing-table
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Routing Tables: Public
       Destinations : 23      Routes : 23

Destination/Mask    Proto  Pre  Cost      Flags NextHop         Interface

      10.0.1.1/32   Direct  0    0          D   127.0.0.1       InLoopBack0
      10.0.4.4/32   OSPF    10   2          D   10.0.111.11 GigabitEthernet0/0/1
    10.0.11.11/32   OSPF    10   1          D   10.0.111.11 GigabitEthernet0/0/1
     10.0.12.0/24   Direct  0    0          D   10.0.12.1       Serial1/0/0
     10.0.12.1/32   Direct  0    0          D   127.0.0.1       InLoopBack0
```

```
     10.0.12.2/32   Direct  0    0           D   10.0.12.2       Serial1/0/0
   10.0.12.255/32   Direct  0    0           D   127.0.0.1       InLoopBack0
     10.0.14.0/24   Direct  0    0           D   10.0.14.1       Serial3/0/0
     10.0.14.1/32   Direct  0    0           D   127.0.0.1       InLoopBack0
     10.0.14.4/32   Direct  0    0           D   10.0.14.4       Serial3/0/0
   10.0.14.255/32   Direct  0    0           D   127.0.0.1       InLoopBack0
     10.0.15.0/24   Direct  0    0           D   10.0.15.1       Serial2/0/0
     10.0.15.1/32   Direct  0    0           D   127.0.0.1       InLoopBack0
     10.0.15.3/32   Direct  0    0           D   10.0.15.3       Serial2/0/0
   10.0.15.255/32   Direct  0    0           D   127.0.0.1       InLoopBack0
    10.0.111.0/24   Direct  0    0           D   10.0.111.1  GigabitEthernet0/0/1
    10.0.111.1/32   Direct  0    0           D   127.0.0.1       InLoopBack0
  10.0.111.255/32   Direct  0    0           D   127.0.0.1       InLoopBack0
    10.0.114.0/24   OSPF    10   2           D   10.0.111.11 GigabitEthernet0/0/1
     127.0.0.0/8    Direct  0    0           D   127.0.0.1       InLoopBack0
     127.0.0.1/32   Direct  0    0           D   127.0.0.1       InLoopBack0
127.255.255.255/32  Direct  0    0           D   127.0.0.1       InLoopBack0
255.255.255.255/32  Direct  0    0           D   127.0.0.1       InLoopBack0

[S1]display ip routing-table
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Routing Tables: Public
       Destinations : 11      Routes : 11

Destination/Mask    Proto  Pre  Cost      Flags NextHop         Interface

      10.0.1.1/32   OSPF   10   1           D   10.0.111.1      Vlanif111
      10.0.4.4/32   OSPF   10   1           D   10.0.114.4      Vlanif114
    10.0.11.11/32   Direct 0    0           D   127.0.0.1       InLoopBack0
    10.0.111.0/24   Direct 0    0           D   10.0.111.11     Vlanif111
   10.0.111.11/32   Direct 0    0           D   127.0.0.1       InLoopBack0
    10.0.114.0/24   Direct 0    0           D   10.0.114.11     Vlanif114
   10.0.114.11/32   Direct 0    0           D   127.0.0.1       InLoopBack0
     10.1.11.0/24   Direct 0    0           D   10.1.11.11      LoopBack1
    10.1.11.11/32   Direct 0    0           D   127.0.0.1       InLoopBack0
     127.0.0.0/8    Direct 0    0           D   127.0.0.1       InLoopBack0
     127.0.0.1/32   Direct 0    0           D   127.0.0.1       InLoopBack0

<R4>display ip routing-table
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Routing Tables: Public
```

```
         Destinations : 18       Routes : 18

Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface

      10.0.1.1/32   OSPF    10   2           D   10.0.114.11 GigabitEthernet0/0/1
      10.0.4.4/32   Direct  0    0           D   127.0.0.1       InLoopBack0
    10.0.11.11/32   OSPF    10   1           D   10.0.114.11 GigabitEthernet0/0/1
     10.0.14.0/24   Direct  0    0           D   10.0.14.4       Serial1/0/0
     10.0.14.1/32   Direct  0    0           D   10.0.14.1       Serial1/0/0
     10.0.14.4/32   Direct  0    0           D   127.0.0.1       InLoopBack0
   10.0.14.255/32   Direct  0    0           D   127.0.0.1       InLoopBack0
     10.0.45.0/24   Direct  0    0           D   10.0.45.4   GigabitEthernet0/0/0
     10.0.45.4/32   Direct  0    0           D   127.0.0.1       InLoopBack0
   10.0.45.255/32   Direct  0    0           D   127.0.0.1       InLoopBack0
    10.0.111.0/24   OSPF    10   2           D   10.0.114.11 GigabitEthernet0/0/1
    10.0.114.0/24   Direct  0    0           D   10.0.114.4  GigabitEthernet0/0/1
    10.0.114.4/32   Direct  0    0           D   127.0.0.1       InLoopBack0
  10.0.114.255/32   Direct  0    0           D   127.0.0.1       InLoopBack0
     127.0.0.0/8    Direct  0    0           D   127.0.0.1       InLoopBack0
     127.0.0.1/32   Direct  0    0           D   127.0.0.1       InLoopBack0
127.255.255.255/32  Direct  0    0           D   127.0.0.1       InLoopBack0
255.255.255.255/32  Direct  0    0           D   127.0.0.1       InLoopBack0
```

在R1、R4、S1上配置BGP，要求使用各自的Loopback 0接口建立对等体连接，配置的时候使用对等体组，组名为**as64512。**

默认情况下，BGP的负载分担是关闭的。在所有的路由器上打开负载分担，设置最大同时使用4条等价路径。

```
[R1]bgp 64512
[R1-bgp]group as64512 internal
[R1-bgp]peer 10.0.11.11 group as64512
[R1-bgp]peer 10.0.11.11 connect-interface LoopBack 0
[R1-bgp]maximum load-balancing 4

[S1]bgp 64512
[S1-bgp]group as64512 internal
[S1-bgp]peer 10.0.4.4 group as64512
[S1-bgp]peer 10.0.4.4 connect-interface LoopBack 0
[S1-bgp]maximum load-balancing 4
[S1-bgp]peer 10.0.1.1 group as64512
[S1-bgp]peer 10.0.1.1 connect-interface LoopBack 0

[R4]bgp 64512
```

```
[R4-bgp]group as64512 internal
[R4-bgp]peer 10.0.11.11 group as64512
[R4-bgp]peer 10.0.11.11 connect-interface LoopBack 0
[R4-bgp]maximum load-balancing 4
```

在R1、R2、R3、R4、R5上配置EBGP，AS的规划如图所示，EBGP全部使用物理接口地址建立对等体关系。

```
[R1]bgp 64512
[R1-bgp]peer 10.0.12.2 as-number 200
[R1-bgp]peer 10.0.15.3 as-number 100

[R2]bgp 200
[R2-bgp]peer 10.0.12.1 as-number 64512
[R2-bgp]peer 10.0.23.3 as-number 100
[R2-bgp]maximum load-balancing 4

[R3]bgp 100
[R3-bgp]peer 10.0.23.2 as-number 200
[R3-bgp]peer 10.0.35.5 as-number 100
[R3-bgp]peer 10.0.15.1 as-number 64512
[R3-bgp]maximum load-balancing 4

[R4]bgp 64512
[R4-bgp]peer 10.0.45.5 as-number 100

[R5]bgp 100
[R5-bgp]peer 10.0.35.3 as-number 100
[R5-bgp]peer 10.0.45.4 as-number 64512
[R5-bgp]maximum load-balancing 4
```

步骤三. 配置 AS-Path 属性

在S1上创建Loopback 1，地址为10.1.11.11/24，使用**network**命令发布到BGP中。

```
[S1]interface loopback 1
[S1-LoopBack1]ip address 10.1.11.11 255.255.255.0
[S1-LoopBack1]bgp 64512
[S1-bgp]network 10.1.11.11 255.255.255.0
```

在R2上观察BGP路由表，可看出这时10.1.11.0/24路由是依据AS-Path属性

来选择下一跳的。

```
[R2]display bgp routing-table

 BGP Local router ID is 10.0.12.2
 Status codes: * - valid, > - best, d - damped,
               h - history,  i - internal, s - suppressed, S - Stale
               Origin : i - IGP, e - EGP, ? - incomplete


 Total Number of Routes: 2
     Network          NextHop       MED       LocPrf    PrefVal Path/Ogn

 *>   10.1.11.0/24   10.0.12.1                            0      64512i
 *                   10.0.23.3                            0      100 64512i
```

由于R1到R4之间的带宽有限，现在希望R2能经由AS100来访问10.1.11.0/24。

这里通过AS-Path来影响选路。

在R1上创建路由策略**as_path**，针对10.1.11.0/24这条路由增加2个重复的AS号。

```
[R1]acl number 2001
[R1-acl-basic-2001]rule 5 permit source 10.1.11.0 0.0.0.255
[R1-acl-basic-2001]route-policy as_path permit node 10
[R1-route-policy]if-match acl 2001
[R1-route-policy]apply as-path 64512 64512 additive
```

然后把该策略应用在R1上，让R2从R1上学习到的这条路由的AS-Path有3个值。

```
[R1]bgp 64512
[R1-bgp]peer 10.0.12.2 route-policy as_path  export
```

在R2上观察BGP路由表。

```
<R2>display bgp routing-table

 BGP Local router ID is 10.0.12.2
 Status codes: * - valid, > - best, d - damped,
               h - history,  i - internal, s - suppressed, S - Stale
               Origin : i - IGP, e - EGP, ? - incomplete
```

```
 Total Number of Routes: 2
     Network            NextHop        MED        LocPrf    PrefVal Path/Ogn

 *>   10.1.11.0/24    10.0.23.3                             0      100 64512i
 *                     10.0.12.1                            0      64512 64512
64512i
```

此时R2经由AS100访问10.1.11.0/24网段。

步骤四. 配置 Origin 属性

观察R3的路由表。

```
<R3>display bgp routing-table

 BGP Local router ID is 10.0.15.3
 Status codes: * - valid, > - best, d - damped,
               h - history,  i - internal, s - suppressed, S - Stale
               Origin : i - IGP, e - EGP, ? - incomplete


 Total Number of Routes: 2
     Network            NextHop        MED       LocPrf    PrefVal Path/Ogn

 *>   10.1.11.0/24     10.0.15.1                            0      64512i
 * i                  10.0.35.5                100         0      64512i
```

到达10.1.11.0/24的下一跳是R1，分析原因。

我们希望R3通过R5访问AS 64512，查看10.1.11.0/24原来的Origin属性是IGP。

这时，我们将R1对R3通告的该路由修改为incomplete。

```
[R1]route-policy 22 permit node 10
[R1-route-policy]if-match acl 2001
[R1-route-policy]apply origin incomplete
[R1-route-policy]bgp 64512
[R1-bgp]peer 10.0.15.3 route-policy 22 export
```

策略生效后观察R3的BGP路由表。

```
<R3>display bgp routing-table
```

```
BGP Local router ID is 10.0.15.3
Status codes: * - valid, > - best, d - damped,
              h - history,  i - internal, s - suppressed, S - Stale
              Origin : i - IGP, e - EGP, ? - incomplete


Total Number of Routes: 2
     Network            NextHop       MED        LocPrf    PrefVal Path/Ogn

*>i  10.1.11.0/24     10.0.35.5                  100         0      64512i
*                     10.0.15.1                              0      64512?
```

这时R3到达网络10.1.11.0/24的下一跳是R5。

步骤五.　配置 Local-Pref 属性

本地优先属性在选路中有很高的优先级。

通过改变本地优先属性可可以影响选路。

在R3上创建Loopback 1，地址为10.1.3.3/24，发布进BGP。

```
[R3]interface loopback 1
[R3-LoopBack1]ip address 10.1.3.3 255.255.255.0
[R3-LoopBack1]bgp 100
[R3-bgp]network 10.1.3.3 255.255.255.0
```

在R5上创建Loopback 1，地址为10.1.5.5/24，发布进BGP。

```
[R5]interface loopback 1
[R5-LoopBack1]ip address 10.1.5.5 255.255.255.0
[R5-LoopBack1]bgp 100
[R5-bgp]network 10.1.5.5 255.255.255.0
```

在S1上观察路由表。

```
[S1]display bgp routing-table

 Total Number of Routes: 5

 BGP Local router ID is 10.0.111.11
 Status codes: * - valid, > - best, d - damped,
               h - history,  i - internal, s - suppressed, S - Stale
               Origin : i - IGP, e - EGP, ? - incomplete
```

```
    Network          NextHop       MED        LocPrf    PrefVal Path/Ogn

*>i  10.1.3.0/24    10.0.1.1       0          100        0       100i
* i                 10.0.4.4                  100        0       100i
*>i  10.1.5.0/24    10.0.1.1                  100        0       100i
* i                 10.0.4.4        0         100        0       100i
*>   10.1.11.0/24   0.0.0.0        0                     0       i
```

现在希望到达网络10.1.5.0/24的流量从R4发送到目标，到达网络10.1.3.0/24的流量从R1发送到目标。

在R4上创建路由策略Pref4，匹配路由10.1.5.0/24，将其本地优先属性修改为110。

R1上创建路由策略Pref1，匹配路由10.1.3.0/24，将其本地优先属性修改为110，然后将策略应用到IBGP的对等体组上。

```
[R4]acl number 2001
[R4-acl-basic-2001]rule 5 permit source 10.1.5.0 0.0.0.255
[R4-acl-basic-2001]quit
[R4]route-policy Pref4 permit node 10
[R4-route-policy]if-match acl 2001
[R4-route-policy]apply local-preference 110
[R4-route-policy]route-policy Pref4 permit node 20
[R4-route-policy]bgp 64512
[R4-bgp]peer as64512 route-policy Pref4 export

[R1]acl number 2002
[R1-acl-basic-2002]rule 5 permit source 10.1.3.0 0.0.0.255
[R1-acl-basic-2002]route-policy Pref1 permit node 10
[R1-route-policy]if-match acl 2002
[R1-route-policy]apply local-preference 110
[R1-route-policy]route-policy Pref1 permit node 20
[R1-route-policy]bgp 64512
[R1-bgp]peer as64512 route-policy Pref1 export
```

在S1上查看BGP路由表。

```
[S1]display bgp routing-table

 Total Number of Routes: 3

 BGP Local router ID is 10.0.111.11
 Status codes: * - valid, > - best, d - damped,
               h - history,  i - internal, s - suppressed, S - Stale
```

```
          Origin : i - IGP, e - EGP, ? - incomplete

     Network            NextHop        MED         LocPrf     PrefVal  Path/Ogn

 *>i  10.1.3.0/24        10.0.1.1        0           110         0       100i
 * i                     10.0.4.4                    100         0       100i
 *>i  10.1.5.0/24        10.0.4.4        0           110         0       100i
 * i                     10.0.1.1         0          100         0       100i
 *>   10.1.11.0/24       0.0.0.0         0                       0       i
```

可以观察到，此时根据Local-Pref属性进行选路，越高越优先。

步骤六. 配置 MED 属性

删除步骤四中通过修改Origin来影响AS100中对10.1.11.0/24选路的路由策略，本实验中通过修改MED值来影响选路。

```
[R1]undo route-policy 22
[R1]bgp 64512
[R1-bgp]undo peer 10.0.15.3 route-policy 22 export
```

在R1上创建路由策略**med**，针对10.1.11.0/24，将MED值修改为100，将该策略应用到对等体R3上。

```
[R1]route-policy med permit node 10
[R1-route-policy]if-match acl 2001
[R1-route-policy]apply cost 100
[R1-route-policy]bgp 64512
[R1-bgp]peer 10.0.15.3 route-policy med export
```

在R3上查看BGP路由表。

```
<R3>dis bgp routing-table

 BGP Local router ID is 10.0.15.3
 Status codes: * - valid, > - best, d - damped,
               h - history,  i - internal, s - suppressed, S - Stale
               Origin : i - IGP, e - EGP, ? - incomplete


 Total Number of Routes: 4
     Network          NextHop        MED        LocPrf    PrefVal  Path/Ogn
```

```
 *>   10.1.3.0/24   0.0.0.0        0                   0      i
 *>i  10.1.5.0/24   10.0.35.5      0         100       0      i
 *>i  10.1.11.0/24  10.0.35.5                100       0      64512i
 *                  10.0.15.1      100                 0      64512i
```

MED值越小越优先。

最后观察现象，可以和步骤四中达到同样的选路效果。

附加实验：思考并验证

思考在完成了步骤六以后，关闭R1的S1/0/0接口，那么在R2上学习到的关于10.1.11.0/24的MED值是多少？

思考可否使用路由策略将AS-Path属性里的某个AS删除？

最终设备配置

```
<R1>display current-configuration
[V200R001C00SPC200]
#
 sysname R1
#
interface Serial1/0/0
 link-protocol ppp
 ip address 10.0.12.1 255.255.255.0
#
interface Serial3/0/0
 link-protocol ppp
 ip address 10.0.14.1 255.255.255.0
#
interface GigabitEthernet0/0/1
 ip address 10.0.111.1 255.255.255.0
#
interface GigabitEthernet0/0/2
 ip address 10.0.15.1 255.255.255.0
#
interface LoopBack0
 ip address 10.0.1.1 255.255.255.255
```

```
#
bgp 64512
 peer 10.0.12.2 as-number 200
 peer 10.0.15.3 as-number 100
 group as64512 internal
 peer 10.0.11.11 as-number 64512
 peer 10.0.11.11 group as64512
 peer 10.0.11.11 connect-interface LoopBack0
 #
 ipv4-family unicast
  undo synchronization
  maximum load-balancing 4
  peer 10.0.12.2 enable
  peer 10.0.12.2 route-policy as_path export
  peer 10.0.15.3 enable
  peer 10.0.15.3 route-policy med export
  peer as64512 enable
  peer as64512 route-policy Pref1 export
  peer 10.0.11.11 enable
  peer 10.0.11.11 group as64512
#
ospf 1
 area 0.0.0.0
  network 10.0.1.1 0.0.0.0
  network 10.0.111.1 0.0.0.0
#
route-policy as_path permit node 10
 if-match acl 2001
 apply as-path 64512 64512 additive
#
route-policy Pref1 permit node 10
 if-match acl 2002
 apply local-preference 110
#
route-policy Pref1 permit node 20
#
route-policy med permit node 10
 if-match acl 2001
 apply cost 100
#
return

<R2>display current-configuration
```

```
[V200R001C00SPC200]
#
 sysname R2
#
interface Serial1/0/0
 link-protocol ppp
 ip address 10.0.12.2 255.255.255.0
#
interface Serial2/0/0
 link-protocol ppp
 ip address 10.0.23.2 255.255.255.0
#
interface LoopBack0
 ip address 10.0.2.2 255.255.255.0
#
bgp 200
 peer 10.0.12.1 as-number 64512
 peer 10.0.23.3 as-number 100
 #
 ipv4-family unicast
  undo synchronization
  maximum load-balancing 4
  peer 10.0.12.1 enable
  peer 10.0.23.3 enable
#
return
```

```
<R3>display current-configuration
[V200R001C00SPC200]
#
 sysname R3
#
interface Serial2/0/0
 link-protocol ppp
 ip address 10.0.23.3 255.255.255.0
#
interface Serial3/0/0
 link-protocol ppp
 ip address 10.0.35.3 255.255.255.0
#
interface GigabitEthernet0/0/2
 ip address 10.0.15.3 255.255.255.0
#
```

```
interface LoopBack0
 ip address 10.0.3.3 255.255.255.255
#
interface LoopBack1
 ip address 10.1.3.3 255.255.255.0
#
bgp 100
 peer 10.0.15.1 as-number 64512
 peer 10.0.23.2 as-number 200
 peer 10.0.35.5 as-number 100
 #
 ipv4-family unicast
  undo synchronization
  network 10.1.3.0 255.255.255.0
  maximum load-balancing 4
  peer 10.0.15.1 enable
  peer 10.0.23.2 enable
  peer 10.0.35.5 enable
#
return

<R4>display current-configuration
[V200R001C00SPC200]
#
 sysname R4
#
interface Serial1/0/0
 link-protocol ppp
 ip address 10.0.14.4 255.255.255.0
#
interface GigabitEthernet0/0/0
 ip address 10.0.45.4 255.255.255.0
#
interface GigabitEthernet0/0/1
 ip address 10.0.114.4 255.255.255.0
#
interface LoopBack0
 ip address 10.0.4.4 255.255.255.255
#
bgp 64512
 peer 10.0.45.5 as-number 100
 group as64512 internal
 peer 10.0.11.11 as-number 64512
```

```
 peer 10.0.11.11 group as64512
 peer 10.0.11.11 connect-interface LoopBack0
 #
 ipv4-family unicast
  undo synchronization
  maximum load-balancing 4
  peer 10.0.45.5 enable
  peer as64512 enable
  peer as64512 route-policy Pref4 export
  peer 10.0.11.11 enable
  peer 10.0.11.11 group as64512
#
ospf 1
 area 0.0.0.0
  network 10.0.114.4 0.0.0.0
  network 10.0.4.4 0.0.0.0
#
route-policy Pref4 permit node 10
 if-match acl 2001
 apply local-preference 110
#
route-policy Pref4 permit node 20
#
return

<R5>display current-configuration
[V200R001C00SPC200]
#
 sysname R5
#
interface Serial1/0/0
 link-protocol ppp
 ip address 10.0.35.5 255.255.255.0
#
interface GigabitEthernet0/0/0
 ip address 10.0.45.5 255.255.255.0
#
interface LoopBack0
 ip address 10.0.5.5 255.255.255.255
#
interface LoopBack1
 ip address 10.1.5.5 255.255.255.0
#
```

```
bgp 100
 peer 10.0.35.3 as-number 100
 peer 10.0.45.4 as-number 64512
 #
 ipv4-family unicast
  undo synchronization
  network 10.1.5.0 255.255.255.0
  maximum load-balancing 4
  peer 10.0.35.3 enable
  peer 10.0.45.4 enable
#
return
```

实验 3-4 BGP 属性与路径选择 2(选做)

学习目的

- 掌握通过修改团体属性来影响路径选择的配置方法
- 掌握使用Route-policy来过滤BGP路由信息的配置方法

拓扑图

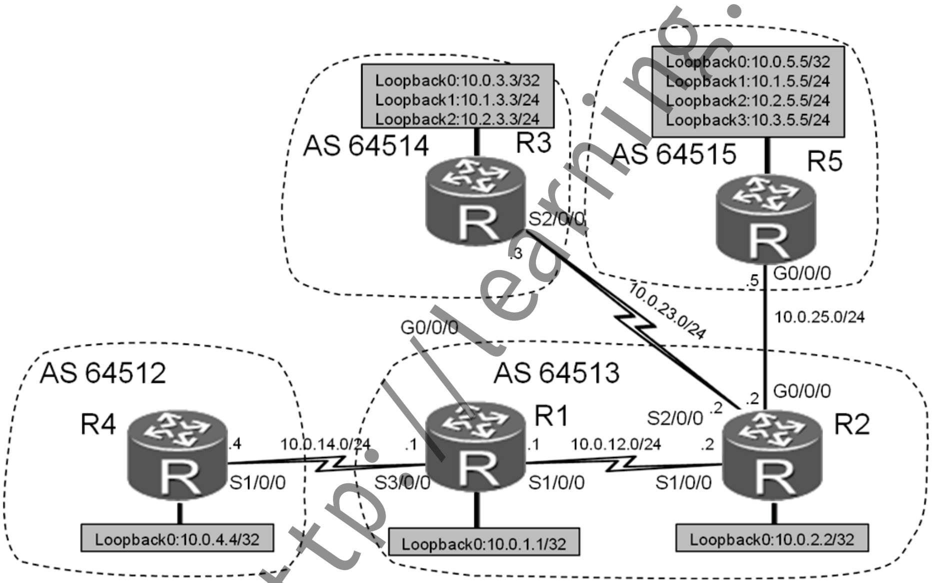

图3-4 BGP属性与路径选择2

场景

你是公司的网络管理员。公司的网络采用了BGP进行互联，BGP的AS号规划如拓扑图中所示。为了公司网络的安全，并非所有分支机构之间的所有部门都能互访。为了控制路由信息的传递，现在你需要使用BGP的团体属性对BGP的路由进行过滤。

学习任务

步骤一.　基础配置与 IP 编址

给所有路由器配置物理接口及Loopback 0接口的IP地址和掩码。注意各Loopback 0接口地址均使用32位掩码。

```
<R1>system-view
Enter system view, return user view with Ctrl+Z.
[R1]interface Serial 1/0/0
[R1-Serial1/0/0]ip address 10.0.12.1 255.255.255.0
[R1-Serial1/0/0]interface Serial 3/0/0
[R1-Serial3/0/0]ip address 10.0.14.1 255.255.255.0
[R1-Serial3/0/0]interface loopback 0
[R1-LoopBack0]ip address 10.0.1.1 255.255.255.255

<R2>system-view
Enter system view, return user view with Ctrl+Z.
[R2]interface Serial 1/0/0
[R2-Serial1/0/0]ip address 10.0.12.2 255.255.255.0
[R2-Serial1/0/0]interface Serial 2/0/0
[R2-Serial2/0/0]ip address 10.0.23.2 255.255.255.0
[R2-Serial2/0/0]interface GigabitEthernet 0/0/0
[R2-GigabitEthernet0/0/0]ip address 10.0.25.2 255.255.255.0
[R2-GigabitEthernet0/0/0]interface loopback 0
[R2-LoopBack0]ip add 10.0.2.2 255.255.255.255

<R3>system-view
Enter system view, return user view with Ctrl+Z.
[R3]interface Serial 2/0/0
[R3-Serial2/0/0]ip address 10.0.23.3 255.255.255.0
[R3-Serial2/0/0]interface loopback 0
[R3-LoopBack0]ip address 10.0.3.3 255.255.255.255

<R4>system-view
Enter system view, return user view with Ctrl+Z.
[R4]interface Serial 1/0/0
[R4-Serial1/0/0]ip address 10.0.14.4 255.255.255.0
[R4-Serial1/0/0]interface loopback 0
[R4-LoopBack0]ip address 10.0.4.4 255.255.255.255

<R5>system-view
```

```
Enter system view, return user view with Ctrl+Z.
[R5]interface GigabitEthernet 0/0/0
[R5-GigabitEthernet0/0/0]ip address 10.0.25.5 255.255.255.0
[R5-GigabitEthernet0/0/0]interface loopback 0
[R5-LoopBack0]ip address 10.0.5.5 255.255.255.255
```

配置完成后，测试直连链路的连通性。

```
<R1>ping -c 1 10.0.12.2
  PING 10.0.12.2: 56  data bytes, press CTRL_C to break
    Reply from 10.0.12.2: bytes=56 Sequence=1 ttl=255 time=40 ms

  --- 10.0.12.2 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 40/40/40 ms

<R1>ping -c 1 10.0.14.4
  PING 10.0.14.4: 56  data bytes, press CTRL_C to break
    Reply from 10.0.14.4: bytes=56 Sequence=1 ttl=255 time=61 ms

  --- 10.0.14.4 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 61/61/61 ms

<R2>ping -c 1 10.0.25.5
  PING 10.0.25.5: 56  data bytes, press CTRL_C to break
    Reply from 10.0.25.5: bytes=56 Sequence=1 ttl=255 time=14 ms

  --- 10.0.25.5 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 14/14/14 ms

<R2>ping -c 1 10.0.23.2
  PING 10.0.23.2: 56  data bytes, press CTRL_C to break
    Reply from 10.0.23.2: bytes=56 Sequence=1 ttl=255 time=2 ms

  --- 10.0.23.2 ping statistics ---
    1 packet(s) transmitted
```

```
 1 packet(s) received
 0.00% packet loss
 round-trip min/avg/max = 2/2/2 ms
```

步骤二. 配置 BGP

R1与R2之间为IBGP邻居关系，其他路由器之间均为EBGP邻居关系，

```
[R1]bgp 64513
[R1-bgp]peer 10.0.12.2 as-number 64513
[R1-bgp]peer 10.0.14.4 as-number 64512

[R2]bgp 64513
[R2-bgp]peer 10.0.12.1 as-number 64513
[R2-bgp]peer 10.0.23.3 as-number 64514
[R2-bgp]peer 10.0.25.5 as-number 64515

[R3]bgp 64514
[R3-bgp]peer 10.0.23.2 as-number 64513

[R4]bgp 64512
[R4-bgp]peer 10.0.14.1 as-number 64513

[R5]bgp 64515
[R5-bgp]peer 10.0.25.2 as-number 64513
```

在配置完BGP之后检查路由器之间的邻居关系建立情况，

```
[R1]display bgp peer

 BGP local router ID : 10.0.12.1
 Local AS number : 64513
 Total number of peers : 2              Peers in established state : 2

  Peer            V          AS  MsgRcvd  MsgSent  OutQ  Up/Down       State  PrefRcv

  10.0.12.2       4       64513        5        6     0  00:03:28  Established       0
  10.0.14.4       4       64512        2        3     0  00:00:39  Established       0

[R2]display bgp peer

 BGP local router ID : 10.0.12.2
 Local AS number : 64513
```

```
 Total number of peers : 3              Peers in established state : 3

  Peer            V          AS  MsgRcvd  MsgSent  OutQ  Up/Down    State     PrefRcv

  10.0.12.1       4        64513        6       5     0   00:04:00   Established    0
  10.0.23.3       4        64514        4       6     0   00:02:44   Established    0
  10.0.25.5       4        64515        2       3     0   00:00:41   Established    0

[R3]display bgp peer

 BGP local router ID : 10.0.23.3
 Local AS number : 64514
 Total number of peers : 1              Peers in established state : 1

  Peer            V           AS  MsgRcvd  MsgSent  OutQ  Up/Down   State     PrefRcv

  10.0.23.2       4        64513        4         4     0   00:02:59  Established   0

 [R4]display bgp peer

 BGP local router ID : 10.0.14.4
 Local AS number : 64512
 Total number of peers : 1              Peers in established state : 1

  Peer            V          AS  MsgRcvd  MsgSent  OutQ   Up/Down   State     PrefRcv

  10.0.14.1       4        64513        3         3    0  00:01:40   Established    0

[R5]display bgp peer

 BGP local router ID : 10.0.25.5
 Local AS number : 64515
 Total number of peers : 1              Peers in established state : 1

  Peer            V           AS  MsgRcvd  MsgSent  OutQ  Up/Down    State     PrefRcv

  10.0.25.2       4        64513        3         3    0   00:01:23   Established    0
```

此时BGP邻居状态全部都是Established状态。

步骤三. 配置普通的团体属性

在R5上创建Loopback1、Loopback2和Loopback3，地址分别为10.1.5.5/24、10.2.5.5/24、10.3.5.5/24，并通过**network**命令发布到BGP中。

```
[R5]interface loopback 1
[R5-LoopBack1]ip address 10.1.5.5 255.255.255.0
[R5-LoopBack1]interface loopback 2
[R5-LoopBack2]ip a  ddress 10.2.5.5 255.255.255.0
[R5-LoopBack2]interface loopback 3
[R5-LoopBack3]ip address 10.3.5.5 255.255.255.0
[R5-LoopBack3]quit
[R5]bgp 64515
[R5-bgp]network 10.1.5.5 255.255.255.0
[R5-bgp]network 10.2.5.5 255.255.255.0
[R5-bgp]network 10.3.5.5 255.255.255.0

[R2]bgp 64513
[R2-bgp]peer 10.0.12.1 next-hop-local
```

在R2和R4上检查该路由信息是否被正确传递。

```
[R2]display bgp routing-table

 BGP Local router ID is 10.0.12.2
 Status codes: * - valid, > - best, d - damped,
               h - history,  i - internal, s - suppressed, S - Stale
               Origin : i - IGP, e - EGP, ? - incomplete


 Total Number of Routes: 3
     Network            NextHop        MED        LocPrf    PrefVal Path/Ogn

 *>   10.1.5.0/24        10.0.25.5      0                     0      64515i
 *>   10.2.5.0/24        10.0.25.5      0                     0      64515i
 *>   10.3.5.0/24        10.0.25.5      0                     0      64515i

[R4]display bgp routing-table

 BGP Local router ID is 10.0.14.4
 Status codes: * - valid, > - best, d - damped,
               h - history,  i - internal, s - suppressed, S - Stale
               Origin : i - IGP, e - EGP, ? - incomplete
```

```
 Total Number of Routes: 3
     Network            NextHop        MED        LocPrf    PrefVal    Path/Ogn

 *>   10.1.5.0/24        10.0.14.1                                  0      64513 64515i
 *>   10.2.5.0/24        10.0.14.1                                  0      64513 64515i
 *>   10.3.5.0/24        10.0.14.1                                  0      64513 64515i
```

在R5上创建路由策略**comm_r5**，对10.1.5.0/24这条路由添加团体属性为100。

```
[R5]acl number 2000
[R5-acl-basic-2000]rule 0 permit source 10.1.5.0 0.0.0.255
[R5-acl-basic-2000]quit
[R5]route-policy comm_r5 permit node 10
[R5-route-policy]if-match acl 2000
[R5-route-policy]apply community 100
[R5-route-policy]quit

[R5]bgp 64515
[R5-bgp]peer 10.0.25.2 route-policy comm_r5 export
```

为了后面实验的需要，需要允许路由器之间通告团体属性，配置所有BGP邻居之间通告Community属性。

```
[R1]bgp 64513
[R1-bgp]peer 10.0.14.4 advertise-community
[R1-bgp]peer 10.0.12.2 advertise-community

[R2]bgp 64513
[R2-bgp]peer 10.0.12.1 advertise-community
[R2-bgp]peer 10.0.23.3 advertise-community
[R2-bgp]peer 10.0.25.5 advertise-community

[R3]bgp 64514
[R3-bgp]peer 10.0.23.2 advertise-community

[R4]bgp 64512
[R4-bgp]peer 10.0.14.1 advertise-community

[R5]bgp 64515
[R5-bgp]peer 10.0.25.2 advertise-community
```

在R2和R4上查看该属性是否被正常传递。

```
<R2>display bgp routing-table community

 BGP Local router ID is 10.0.12.2
 Status codes: * - valid, > - best, d - damped,
               h - history,  i - internal, s - suppressed, S - Stale
               Origin : i - IGP, e - EGP, ? - incomplete


 Total Number of Routes: 5
     Network            NextHop        MED        LocPrf    PrefVal Community

*>   10.1.5.0/24        10.0.25.5       0                    0       <0:100>
```

```
<R4>display bgp routing-table community

 BGP Local router ID is 10.0.12.2
 Status codes: * - valid, > - best, d - damped,
               h - history,  i - internal, s - suppressed, S - Stale
               Origin : i - IGP, e - EGP, ? - incomplete


 Total Number of Routes: 5
     Network            NextHop        MED        LocPrf    PrefVal Community

*>   10.1.5.0/24        10.0.25.5       0                    0       <0:100>
```

步骤四. 配置特殊的团体属性值

在R5上使用路由策略为路由10.2.5.0/24添加特殊的团体属性no-export ,使用路由策略为10.3.5.0/24添加特殊的团体属性no-advertise。

这时只需要在R5已创建的路由策略**comm_r5**的基础上添加入两个新节点语句即可。

```
[R5]acl 2001
[R5-acl-basic-2001]rule 0 permit source 10.2.5.0 0.0.0.255
[R5-acl-basic-2001]quit

[R5]route-policy comm_r5 permit node 20
[R5-route-policy]if-match acl 2001
[R5-route-policy]apply community no-export
```

```
[R5]acl number 2002
[R5-acl-basic-2002]rule 0 permit source 10.3.5.0 0.0.0.255
[R5-acl-basic-2002]quit

[R5]route-policy comm_r5 permit node 30
[R5-route-policy]if-match acl 2002
[R5-route-policy]apply community no-advertise
```

在R2上使用查看学习到的路由的团体属性信息。

```
<R2>dis bgp routing-table community

 BGP Local router ID is 10.0.12.2
 Status codes: * - valid, > - best, d - damped,
               h - history,  i - internal, s - suppressed, S - Stale
               Origin : i - IGP, e - EGP, ? - incomplete


 Total Number of Routes: 4
     Network          NextHop        MED        LocPrf    PrefVal Community

 *>   10.1.5.0/24       10.0.25.5      0                     0      <0:100>
 *>   10.2.5.0/24       10.0.25.5      0                     0      no-export
 *>   10.3.5.0/24       10.0.25.5      0                     0      no-advertise
```

在R2、R1、R4上分别观察BGP路由表观察这几条路由的传递情况。

```
<R2>display bgp routing-table

 BGP Local router ID is 10.0.12.2
 Status codes: * - valid, > - best, d - damped,
               h - history,  i - internal, s - suppressed, S - Stale
               Origin : i - IGP, e - EGP, ? - incomplete


 Total Number of Routes: 3
     Network          NextHop        MED        LocPrf    PrefVal Path/Ogn

 *>   10.1.5.0/24       10.0.25.5      0                     0      64515i
 *>   10.2.5.0/24       10.0.25.5      0                     0      64515i
 *>   10.3.5.0/24       10.0.25.5      0                     0      64515i

<R1>display bgp routing-table
```

```
 BGP Local router ID is 10.0.12.1
 Status codes: * - valid, > - best, d - damped,
               h - history,  i - internal, s - suppressed, S - Stale
               Origin : i - IGP, e - EGP, ? - incomplete


 Total Number of Routes: 2
      Network            NextHop        MED        LocPrf    PrefVal Path/Ogn

 *>i  10.1.5.0/24         10.0.12.2      0          100        0      64515i
 *>i  10.2.5.0/24         10.0.12.2      0          100        0      64515i
[R4]display bgp routing-table

 BGP Local router ID is 10.0.14.4
 Status codes: * - valid, > - best, d - damped,
               h - history,  i - internal, s - suppressed, S - Stale
               Origin : i - IGP, e - EGP, ? - incomplete


 Total Number of Routes: 1
      Network            NextHop        MED        LocPrf    PrefVal    Path/Ogn

 *>   10.1.5.0/24        10.0.14.1                             0      64513 64515i
```

对于使用了特殊团体属性no-export后的BGP路由10.2.5.0/24，R2不向其AS外发送，但会发布给其AS内的R1。而对于使用no-advertise属性的BGP路由10.3.5.0/24，R2则不向任何对等体发送。

步骤五.　配置团体属性在地址汇总中的应用

在R3上创建Loopback 1和Loopback2，地址分别配置为10.1.3.3/24和10.2.3.3/24，并通过**network**命令发布到BGP中。

```
[R3]interface LoopBack 1
[R3-LoopBack1]ip address 10.1.3.3 255.255.255.0
[R3-LoopBack1]interface loopback 2
[R3-LoopBack2]ip address 10.2.3.3 255.255.255.0
[R3-LoopBack2]quit
[R3]bgp 64514
[R3-bgp]network 10.1.3.3 255.255.255.0
```

```
[R3-bgp]network 10.2.3.3 255.255.255.0
```

现在有一个需求，我们需要将R5发布的10.1.5.0/24和R3发布的10.2.3.0/24汇总成一个A类网段10.0.0.0/8。通告时抑制明细路由，并且该汇总路由最后通告给R4时携带的团体属性为200。对路由10.1.3.0/24保留明细通告给R4。

为了实现这个需求，我们在R3上创建一个名为**comm_r3**的路由策略，对R3发布的路由10.2.3.0/24加上100的团体属性。

```
[R3]acl number 2001
[R3-acl-basic-2001]rule 0 permit source 10.2.3.0 0.0.0.255
[R3-acl-basic-2001]route-policy comm_r3 permit node 10
[R3-route-policy]if-match acl 2001
[R3-route-policy]apply community 100
[R3-route-policy]route-policy comm_r3 permit node 20
[R3-route-policy]quit
[R3]bgp 64514
[R3-bgp]peer 10.0.23.2 route-policy comm_r3 export
```

在R1上观察学习到的10.1.5.0/24和10.2.3.0/24是否都带上了100的团体属性。

```
<R1>display bgp routing-table community

 BGP Local router ID is 10.0.12.1
 Status codes: * - valid, > - best, d - damped,
               h - history,  i - internal, s - suppressed, S - Stale
               Origin : i - IGP, e - EGP, ? - incomplete


 Total Number of Routes: 3
     Network          NextHop        MED        LocPrf    PrefVal Community

 *>i  10.1.5.0/24        10.0.12.2       0          100        0       <0:100>
 *>i  10.2.3.0/24        10.0.12.2       0          100        0       <0:100>
 *>i  10.2.5.0/24        10.0.12.2       0          100        0       no-export
```

创建一个团体属性过滤列表，过滤出团体属性为100的路由。

```
[R1]ip community-filter 1 permit 100
```

创建一个名为**match_comm**的路由策略，匹配团体属性为100的路由。

```
[R1]route-policy match_comm permit node 10
[R1-route-policy]if-match community-filter 1
```

创建一个名为**add_comm**的路由策略，为汇总后路由添加团体属性为200:1。

```
[R1]route-policy add_comm permit node 10
[R1-route-policy]apply community 200:1 additive
```

在R1上进行地址汇总，定义对匹配策略**match_comm**的路由进行汇总，并使用策略**add_comm**添加团体属性。

```
[R1]bgp 64513
[R1-bgp]aggregate 10.0.0.0 255.0.0.0 detail-suppressed origin-policy match_comm
attribute-policy add_comm
```

在R4上查看BGP路由表。

```
<R4>display bgp routing-table

 BGP Local router ID is 10.0.14.4
 Status codes: * - valid, > - best, d - damped,
               h - history,  i - internal, s - suppressed, S - Stale
               Origin : i - IGP, e - EGP, ? - incomplete


 Total Number of Routes: 2
     Network          NextHop        MED        LocPrf    PrefVal Path/Ogn

 *>  10.0.0.0          10.0.14.1                            0      64513i
 *>  10.1.3.0/24       10.0.14.1                            0      64513 64514i
```

在R4上观察汇总路由的团体属性。

```
<R4>display bgp routing-table community

 BGP Local router ID is 10.0.14.4
 Status codes: * - valid, > - best, d - damped,
               h - history,  i - internal, s - suppressed, S - Stale
               Origin : i - IGP, e - EGP, ? - incomplete


 Total Number of Routes: 1
     Network          NextHop        MED        LocPrf    PrefVal Community

 *>  10.0.0.0          10.0.14.1                              0        <200:1>
```

附加实验：思考并验证

在步骤四中，如果将10.2.5.0/24这条路由的属性也改成no advertise，思考此时在R2、R1、R4上再分别查看BGP路由表，这几条路由的传递情况。

思考如何在R4上实现同时保留10.1.3.0/24和10.2.3.0/24这两条路由的明细，仅抑制路由10.1.5.0/24的明细。

最终设备配置

```
<R1>display current-configuration
[V200R001C00SPC200]
#
 sysname R1
#
interface Serial1/0/0
 link-protocol ppp
 ip address 10.0.12.1 255.255.255.0
#
interface Serial3/0/0
 link-protocol ppp
 ip address 10.0.14.1 255.255.255.0
#
interface LoopBack0
 ip address 10.0.1.1 255.255.255.255
#
bgp 64513
 peer 10.0.12.2 as-number 64513
 peer 10.0.14.4 as-number 64512
 #
 ipv4-family unicast
  undo synchronization
  aggregate 10.0.0.0 255.0.0.0 detail-suppressed origin-policy match_comm
attribute-policy add_comm
  peer 10.0.12.2 enable
  peer 10.0.12.2 advertise-community
  peer 10.0.14.4 enable
  peer 10.0.14.4 advertise-community
#
```

```
route-policy match_comm permit node 10
 if-match community-filter 1
#
route-policy add_comm permit node 10
 apply community 200:1 additive
#
 ip community-filter 1 permit 100
#
return

<R2>display current-configuration
[V200R001C00SPC200]
#
 sysname R2
#
interface Serial1/0/0
 link-protocol ppp
 ip address 10.0.12.2 255.255.255.0
#
interface Serial2/0/0
 link-protocol ppp
 ip address 10.0.23.2 255.255.255.0
#
interface GigabitEthernet0/0/0
 ip address 10.0.25.2 255.255.255.0
#
interface LoopBack0
 ip address 10.0.2.2 255.255.255.255
#
bgp 64513
 peer 10.0.12.1 as-number 64513
 peer 10.0.23.3 as-number 64514
 peer 10.0.25.5 as-number 64515
 #
 ipv4-family unicast
  undo synchronization
  peer 10.0.12.1 enable
  peer 10.0.12.1 next-hop-local
  peer 10.0.12.1 advertise-community
  peer 10.0.23.3 enable
  peer 10.0.23.3 advertise-community
  peer 10.0.25.5 enable
  peer 10.0.25.5 advertise-community
```

```
#
return

<R3>display current-configuration
[V200R001C00SPC200]
#
 sysname R3
#
acl number 2001
 rule 0 permit source 10.2.3.0 0.0.0.255
#
interface Serial2/0/0
 link-protocol ppp
 ip address 10.0.23.3 255.255.255.0
#
interface LoopBack0
 ip address 10.0.3.3 255.255.255.255
#
interface LoopBack1
 ip address 10.1.3.3 255.255.255.0
#
interface LoopBack2
 ip address 10.2.3.3 255.255.255.0
#
bgp 64514
 peer 10.0.23.2 as-number 64513
 #
 ipv4-family unicast
  undo synchronization
  network 10.1.3.0 255.255.255.0
  network 10.2.3.0 255.255.255.0
  peer 10.0.23.2 enable
  peer 10.0.23.2 route-policy comm_r3 export
  peer 10.0.23.2 advertise-community
#
route-policy comm_r3 permit node 10
 if-match acl 2001
 apply community 100
#
route-policy comm_r3 permit node 20
#
return
```

```
<R4>display current-configuration
[V200R001C00SPC200]
#
 sysname R4
#
interface Serial1/0/0
 link-protocol ppp
 ip address 10.0.14.4 255.255.255.0
#
interface LoopBack0
 ip address 10.0.4.4 255.255.255.255
#
bgp 64512
 peer 10.0.14.1 as-number 64513
 #
 ipv4-family unicast
  undo synchronization
  peer 10.0.14.1 enable
  peer 10.0.14.1 advertise-community
#
Return

<R5>display current-configuration
[V200R001C00SPC200]
#
 sysname R5
#
interface GigabitEthernet0/0/0
 ip address 10.0.25.5 255.255.255.0
#
interface LoopBack0
 ip address 10.0.5.5 255.255.255.255
#
interface LoopBack1
 ip address 10.1.5.5 255.255.255.0
#
interface LoopBack2
 ip address 10.2.5.5 255.255.255.0
#
interface LoopBack3
 ip address 10.3.5.5 255.255.255.0
#
bgp 64515
```

```
 peer 10.0.25.2 as-number 64513
 #
 ipv4-family unicast
  undo synchronization
  network 10.1.5.0 255.255.255.0
  network 10.2.5.0 255.255.255.0
  network 10.3.5.0 255.255.255.0
  peer 10.0.25.2 enable
  peer 10.0.25.2 route-policy comm_r5 export
  peer 10.0.25.2 advertise-community
#
route-policy comm_r5 permit node 10
 if-match acl 2000
 apply community 100
#
route-policy comm_r5 permit node 20
 if-match acl 2001
 apply community no-export
#
route-policy comm_r5 permit node 30
 if-match acl 2002
 apply community no-advertise
#
return
```

实验 3-5 BGP 多宿主

学习目的

- 掌握BGP多宿主环境中仅使用缺省路由时的配置方法
- 掌握BGP多宿主环境中使用缺省路由过滤部分路由的配置方法
- 掌握BGP多宿主环境中，仅使用BGP路由时的配置方法

拓扑图

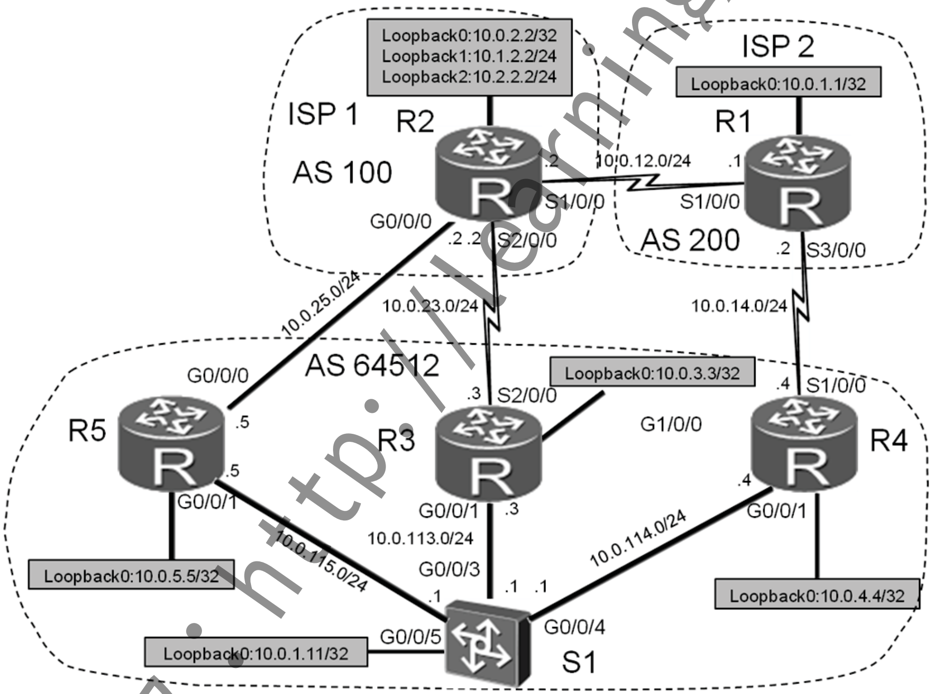

图3-5 BGP多宿主

场景

你是公司的网络管理员。公司的网络采用了BGP协议接入了运营商ISP1。公司自己采用了私有的AS号64512，ISP1的AS号为100，公司从两台路由器分别接入ISP1。起初公司采用默认路由的方式通过运营商接入Internet，随着公司的

发展，这种默认路由的接入方式已不能满足选路的需求，你需要把一部分Internet路由引入到公司的AS中。一段时间以后，公司又租用了一条线路接入到ISP2，ISP2的AS号为200。最终公司实现了由BGP选路的多宿主网络。

学习任务

步骤一. 基础配置与 IP 编址

给所有路由器配置物理接口及Loopback接口的IP地址和掩码。注意各Loopback 0接口均使用32位掩码。

```
<R1>system-view
Enter system view, return user view with Ctrl+Z.
[R1]interface Serial 3/0/0
[R1-Serial3/0/0]ip address 10.0.14.1 255.255.255.0
[R1-Serial3/0/0]interface loopback 0
[R1-LoopBack0]ip address 10.0.1.1 255.255.255.255

<R2>system-view
Enter system view, return user view with Ctrl+Z.
[R2]interface Serial 1/0/0
[R2-Serial1/0/0]ip address 10.0.12.2 255.255.255.0
[R2-Serial1/0/0]interface Serial 2/0/0
[R2-Serial2/0/0]ip address 10.0.23.2 255.255.255.0
[R2-Serial2/0/0]interface GigabitEthernet 0/0/0
[R2-GigabitEthernet0/0/0]ip address 10.0.25.2 255.255.255.0
[R2-GigabitEthernet0/0/0]interface loopback 0
[R2-LoopBack0]ip address 10.0.2.2 255.255.255.255

<R3>system-view
Enter system view, return user view with Ctrl+Z.
[R3]interface Serial 2/0/0
[R3-Serial2/0/0]ip address 10.0.23.3 255.255.255.0
[R3-Serial2/0/0]interface GigabitEthernet 0/0/1
[R3-GigabitEthernet0/0/1]ip address 10.0.113.3 255.255.255.0
[R3-GigabitEthernet0/0/1]interface loopback 0
[R3-LoopBack0]ip address 10.0.3.3 255.255.255.255

<R4>system-view
Enter system view, return user view with Ctrl+Z.
[R4]interface Serial 1/0/0
[R4-Serial1/0/0]ip address 10.0.14.4 255.255.255.0
```

```
[R4-Serial1/0/0]interface GigabitEthernet 0/0/1
[R4-GigabitEthernet0/0/1]ip address 10.0.114.4 255.255.255.0
[R4-GigabitEthernet0/0/1]interface LoopBack 0
[R4-LoopBack0]ip address 10.0.4.4 255.255.255.255

<R5>system-view
Enter system view, return user view with Ctrl+Z.
[R5]interface GigabitEthernet 0/0/0
[R5-GigabitEthernet0/0/0]ip address 10.0.25.5 255.255.255.0
[R5-GigabitEthernet0/0/0]interface GigabitEthernet 0/0/1
[R5-GigabitEthernet0/0/1]ip address 10.0.115.5 255.255.255.0
[R5-GigabitEthernet0/0/1]interface loopback 0
[R5-LoopBack0]ip address 10.0.5.5 255.255.255.255
```

配置完成后，测试直连链路的连通性。

```
<R1>ping -c 1 10.0.14.4
  PING 10.0.14.4: 56  data bytes, press CTRL_C to break
    Reply from 10.0.14.4: bytes=56 Sequence=1 ttl=255 time=33 ms

  --- 10.0.14.4 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 33/33/33 ms

<R1>ping -c 1 10.0.12.2
  PING 10.0.12.2: 56  data bytes, press CTRL_C to break
    Reply from 10.0.12.2: bytes=56 Sequence=1 ttl=255 time=34 ms

  --- 10.0.12.2 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 34/34/34 ms

<R2>ping -c 1 10.0.25.5
  PING 10.0.25.5: 56  data bytes, press CTRL_C to break
    Reply from 10.0.25.5: bytes=56 Sequence=1 ttl=255 time=13 ms

  --- 10.0.25.5 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
```

```
    round-trip min/avg/max = 13/13/13 ms

<R2>ping -c 1 10.0.23.3
  PING 10.0.23.3: 56  data bytes, press CTRL_C to break
    Reply from 10.0.23.3: bytes=56 Sequence=1 ttl=255 time=39 ms

  --- 10.0.23.3 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 39/39/39 ms
```

步骤二.　配置 IGP 及 BGP

在AS 64512内部使用OSPF作为IGP，所有设备属于区域0。

R3的G0/0/1和Loopback 0连接的运行OSPF。

```
[R3]ospf
[R3-ospf-1]area 0
[R3-ospf-1-area-0.0.0.0]network 10.0.113.3 0.0.0.0
[R3-ospf-1-area-0.0.0.0]network 10.0.3.3 0.0.0.0
```

R4的G0/0/1和Loopback 0连接的网段运行OSPF。

```
[R4]ospf
[R4-ospf-1]area 0
[R4-ospf-1-area-0.0.0.0]network 10.0.114.4 0.0.0.0
[R4-ospf-1-area-0.0.0.0]network 10.0.4.4 0.0.0.0
```

R5的G0/0/1和Loopback 0连接的网段运行OSPF，

```
[R5]ospf
[R5-ospf-1]area 0
[R5-ospf-1-area-0.0.0.0]network 10.0.115.5 0.0.0.0
[R5-ospf-1-area-0.0.0.0]network 10.0.5.5 0.0.0.0
```

在S1上创建Vlan13，配置Vlanif地址与R3进行互联。

创建Vlan14，配置Vlanif地址与R4进行互联。

创建Vlan15，配置Vlanif地址与R5进行互联。

互联接口使用Access模式，Vlanif 13、Vlanif 14 、Vlanif 15和Loopback 0连接的网段运行OSPF。

```
[S1]vlan batch 13 to 15
[S1]interface vlan 13
[S1-Vlanif13]ip address 10.0.113.1 255.255.255.0
[S1-Vlanif13]interface vlan 14
[S1-Vlanif14]ip address 10.0.114.1 255.255.255.0
[S1-Vlanif14]interface vlan 15
[S1-Vlanif15]ip address 10.0.115.1 255.255.255.0
[S1-Vlanif15]interface g0/0/3
[S1-GigabitEthernet0/0/3]port link-type access
[S1-GigabitEthernet0/0/3]port default vlan 13
[S1-GigabitEthernet0/0/3]interface g0/0/4
[S1-GigabitEthernet0/0/4]port link-type access
[S1-GigabitEthernet0/0/4]port default vlan 14
[S1-GigabitEthernet0/0/4]interface g0/0/5
[S1-GigabitEthernet0/0/5]port link-type access
[S1-GigabitEthernet0/0/5]port default vlan 15
[S1-GigabitEthernet0/0/5]interface loopback 0
[S1-LoopBack0]ip add 10.0.1.11 255.255.255
[S1-LoopBack0]ospf
[S1-ospf-1]area 0
[S1-ospf-1-area-0.0.0.0]network 10.0.113.1 0.0.0.0
[S1-ospf-1-area-0.0.0.0]network 10.0.114.1 0.0.0.0
[S1-ospf-1-area-0.0.0.0]network 10.0.115.1 0.0.0.0
[S1-ospf-1-area-0.0.0.0]network 10.0.1.11 0.0.0.0
```

检查是否学习到其他设备的Loopback 0接口连接网段的路由。

```
<R3>display ip routing-table
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Routing Tables: Public
        Destinations : 17       Routes : 17

Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface

     10.0.1.11/32   OSPF    10   1           D   10.0.113.1  GigabitEthernet0/0/1
      10.0.3.3/32   Direct  0    0           D   127.0.0.1       InLoopBack0
      10.0.4.4/32   OSPF    10   2           D   10.0.113.1  GigabitEthernet0/0/1
      10.0.5.5/32   OSPF    10   2           D   10.0.113.1  GigabitEthernet0/0/1
     10.0.23.0/24   Direct  0    0           D   10.0.23.3       Serial2/0/0
     10.0.23.2/32   Direct  0    0           D   10.0.23.2       Serial2/0/0
     10.0.23.3/32   Direct  0    0           D   127.0.0.1       InLoopBack0
   10.0.23.255/32   Direct  0    0           D   127.0.0.1       InLoopBack0
    10.0.113.0/24   Direct  0    0           D   10.0.113.3  GigabitEthernet0/0/1
```

```
    10.0.113.3/32  Direct  0    0           D  127.0.0.1       InLoopBack0
  10.0.113.255/32  Direct  0    0           D  127.0.0.1       InLoopBack0
    10.0.114.0/24  OSPF    10   2           D  10.0.113.1  GigabitEthernet0/0/1
    10.0.115.0/24  OSPF    10   2           D  10.0.113.1  GigabitEthernet0/0/1
      127.0.0.0/8  Direct  0    0           D  127.0.0.1       InLoopBack0
     127.0.0.1/32  Direct  0    0           D  127.0.0.1       InLoopBack0
127.255.255.255/32 Direct  0    0           D  127.0.0.1       InLoopBack0
255.255.255.255/32 Direct  0    0           D  127.0.0.1       InLoopBack0

<R4>display ip routing-table
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Routing Tables: Public
        Destinations : 17      Routes : 17

Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface

      10.0.1.11/32  OSPF    10   1           D  10.0.114.1  GigabitEthernet0/0/1
       10.0.3.3/32  OSPF    10   2           D  10.0.114.1  GigabitEthernet0/0/1
       10.0.4.4/32  Direct  0    0           D  127.0.0.1       InLoopBack0
       10.0.5.5/32  OSPF    10   2           D  10.0.114.1  GigabitEthernet0/0/1
      10.0.14.0/24  Direct  0    0           D  10.0.14.4       Serial1/0/0
      10.0.14.1/32  Direct  0    0           D  10.0.14.1       Serial1/0/0
      10.0.14.4/32  Direct  0    0           D  127.0.0.1       InLoopBack0
    10.0.14.255/32  Direct  0    0           D  127.0.0.1       InLoopBack0
     10.0.113.0/24  OSPF    10   2           D  10.0.114.1  GigabitEthernet0/0/1
     10.0.114.0/24  Direct  0    0           D  10.0.114.4  GigabitEthernet0/0/1
     10.0.114.4/32  Direct  0    0           D  127.0.0.1       InLoopBack0
   10.0.114.255/32  Direct  0    0           D  127.0.0.1       InLoopBack0
     10.0.115.0/24  OSPF    10   2           D  10.0.114.1  GigabitEthernet0/0/1
      127.0.0.0/8   Direct  0    0           D  127.0.0.1       InLoopBack0
      127.0.0.1/32  Direct  0    0           D  127.0.0.1       InLoopBack0
127.255.255.255/32  Direct  0    0           D  127.0.0.1       InLoopBack0
255.255.255.255/32  Direct  0    0           D  127.0.0.1       InLoopBack0

<R5>display ip routing-table
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Routing Tables: Public
        Destinations : 16      Routes : 16

Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface
```

```
     10.0.1.11/32   OSPF    10   1           D   10.0.115.1 GigabitEthernet0/0/1
      10.0.3.3/32   OSPF    10   2           D   10.0.115.1  GigabitEthernet0/0/1
      10.0.4.4/32   OSPF    10   2           D   10.0.115.1  GigabitEthernet0/0/1
      10.0.5.5/32   Direct  0    0           D   127.0.0.1       InLoopBack0
     10.0.25.0/24   Direct  0    0           D   10.0.25.5   GigabitEthernet0/0/0
     10.0.25.5/32   Direct  0    0           D   127.0.0.1       InLoopBack0
   10.0.25.255/32   Direct  0    0           D   127.0.0.1       InLoopBack0
    10.0.113.0/24   OSPF    10   2           D   10.0.115.1  GigabitEthernet0/0/1
    10.0.114.0/24   OSPF    10   2           D   10.0.115.1  GigabitEthernet0/0/1
    10.0.115.0/24   Direct  0    0           D   10.0.115.5  GigabitEthernet0/0/1
    10.0.115.5/32   Direct  0    0           D   127.0.0.1       InLoopBack0
  10.0.115.255/32   Direct  0    0           D   127.0.0.1       InLoopBack0
     127.0.0.0/8    Direct  0    0           D   127.0.0.1       InLoopBack0
     127.0.0.1/32   Direct  0    0           D   127.0.0.1       InLoopBack0
127.255.255.255/32  Direct  0    0           D   127.0.0.1       InLoopBack0
255.255.255.255/32  Direct  0    0           D   127.0.0.1       InLoopBack0

[S1]display ip routing-table
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Routing Tables: Public
       Destinations : 12      Routes : 12

Destination/Mask    Proto  Pre  Cost      Flags NextHop         Interface


     10.0.1.11/32   Direct 0    0           D   127.0.0.1       InLoopBack0
      10.0.3.3/32   OSPF   10   1           D   10.0.113.3      Vlanif13
      10.0.4.4/32   OSPF   10   1           D   10.0.114.4      Vlanif14
      10.0.5.5/32   OSPF   10   1           D   10.0.115.5      Vlanif15
    10.0.113.0/24   Direct 0    0           D   10.0.113.1      Vlanif13
    10.0.113.1/32   Direct 0    0           D   127.0.0.1       InLoopBack0
    10.0.114.0/24   Direct 0    0           D   10.0.114.1      Vlanif14
    10.0.114.1/32   Direct 0    0           D   127.0.0.1       InLoopBack0
    10.0.115.0/24   Direct 0    0           D   10.0.115.1      Vlanif15
    10.0.115.1/32   Direct 0    0           D   127.0.0.1       InLoopBack0
     127.0.0.0/8    Direct 0    0           D   127.0.0.1       InLoopBack0
     127.0.0.1/32   Direct 0    0           D   127.0.0.1       InLoopBack0
```

在R2、R3、R5上配置EBGP。

AS的规划如图所示，EBGP全部使用物理接口地址建立对等体关系。S1不运行BGP。

```
[R2]bgp 100
```

```
[R2-bgp]peer 10.0.25.5 as-number 64512
[R2-bgp]peer 10.0.23.3 as-number 64512

[R3]bgp 64512
[R3-bgp]peer 10.0.23.2 as-number 100

[R5]bgp 64512
[R5-bgp]peer 10.0.25.2 as-number 100
```

完成后检查BGP邻居是否完全建立。

```
[R2]display bgp peer

 BGP local router ID : 10.0.12.2
 Local AS number : 100
 Total number of peers : 2              Peers in established state : 2

  Peer            V         AS  MsgRcvd  MsgSent  OutQ  Up/Down  State      PrefRcv

  10.0.23.3        4       64512        7        9     0 00:05:55 Established       0
  10.0.25.5        4       64512        6        7     0 00:04:17 Established       0

[R3]display bgp peer

 BGP local router ID : 10.0.23.3
 Local AS number : 64512
 Total number of peers : 1              Peers in established state : 1

  Peer            V         AS  MsgRcvd  MsgSent  OutQ  Up/Down   State     PrefRcv

  10.0.23.2        4        100        8        8     0 00:06:09 Established       0

<R5>display bgp peer

 BGP local router ID : 10.0.25.5
 Local AS number : 64512
 Total number of peers : 1              Peers in established state : 1

  Peer            V         AS  MsgRcvd  MsgSent  OutQ  Up/Down   State     PrefRcv

  10.0.25.2        4        100        7        7     0  00:05:31 Established       0
```

步骤三. 配置仅使用缺省路由归属到单个 ISP

默认情况下，BGP的负载分担是关闭的。在所有的路由器上打开负载分担，设置为最大4条路径。

```
[R1]bgp 200
[R1-bgp]maximum load-balancing 4

[R2]bgp 100
[R2-bgp]maximum load-balancing 4

[R3]bgp 64512
[R3-bgp]maximum load-balancing 4

[R4]bgp 64512
[R4-bgp]maximum load-balancing 4

[R5]bgp 64512
[R5-bgp]maximum load-balancing 4
```

在R2上创建Loopback 1和Loopback 2，地址分别为10.1.2.2/24和10.2.2.2/24。使用**network**命令将这两个网段发布到BGP中，

```
[R2]interface LoopBack 1
[R2-LoopBack1]ip address 10.1.2.2 255.255.255.0
[R2]interface LoopBack 2
[R2-LoopBack2]ip address 10.2.2.2 255.255.255.0

[R2]bgp 100
[R2-bgp]network 10.1.2.2 255.255.255.0
[R2-bgp]network 10.2.2.0 255.255.255.0
```

在R3和R5上分别检查是否学习到该路由。

```
[R3]display bgp routing-table

 BGP Local router ID is 10.0.23.3
 Status codes: * - valid, > - best, d - damped,
               h - history,  i - internal, s - suppressed, S - Stale
               Origin : i - IGP, e - EGP, ? - incomplete


 Total Number of Routes: 2
```

```
     Network            NextHop        MED        LocPrf    PrefVal Path/Ogn

 *>   10.1.2.0/24        10.0.23.2      0                     0      100i
 *>   10.2.2.0/24        10.0.23.2      0                     0      100i

<R5>display bgp  routing-table

 BGP Local router ID is 10.0.25.5
 Status codes: * - valid, > - best, d - damped,
               h - history,  i - internal, s - suppressed, S - Stale
               Origin : i - IGP, e - EGP, ? - incomplete


 Total Number of Routes: 2
     Network            NextHop        MED        LocPrf    PrefVal Path/Ogn

 *>   10.1.2.0/24        10.0.25.2      0                     0      100i
 *>   10.2.2.0/24        10.0.25.2      0                     0      100i
```

在这里，R3接入到ISP1的线路为主用线路，R5接入到ISP1的线路为备用线路。在R3和R5上分别将OSPF的路由通过**import-route**命令引入进BGP，

```
[R3]bgp 64512
[R3-bgp]import-route ospf 1

[R5]bgp 64512
[R5-bgp]import-route ospf 1
```

在R3和R5上向区域0内强制发布默认路由，并采用第一类外部路由类型发布。设置R3发布的默认路由cost为20，R5发布的默认路由cost为40。

```
[R3]ospf
[R3-ospf-1]default-route-advertise always cost 20 type 1

[R5]ospf
[R5-ospf-1]default-route-advertise always cost 40 type 1
```

在S1上查看路由表。

```
[S1]display ip routing-table
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Routing Tables: Public
         Destinations : 13       Routes : 13
```

```
Destination/Mask    Proto  Pre  Cost      Flags NextHop         Interface

        0.0.0.0/0   O_ASE  150  21          D   10.0.113.3      Vlanif13
      10.0.1.11/32  Direct 0    0           D   127.0.0.1       InLoopBack0
       10.0.3.3/32  OSPF   10   1           D   10.0.113.3      Vlanif13
       10.0.4.4/32  OSPF   10   1           D   10.0.114.4      Vlanif14
       10.0.5.5/32  OSPF   10   1           D   10.0.115.5      Vlanif15
     10.0.113.0/24  Direct 0    0           D   10.0.113.1      Vlanif13
     10.0.113.1/32  Direct 0    0           D   127.0.0.1       InLoopBack0
     10.0.114.0/24  Direct 0    0           D   10.0.114.1      Vlanif14
     10.0.114.1/32  Direct 0    0           D   127.0.0.1       InLoopBack0
     10.0.115.0/24  Direct 0    0           D   10.0.115.1      Vlanif15
     10.0.115.1/32  Direct 0    0           D   127.0.0.1       InLoopBack0
      127.0.0.0/8   Direct 0    0           D   127.0.0.1       InLoopBack0
      127.0.0.1/32  Direct 0    0           D   127.0.0.1       InLoopBack0
```

在S1上查看到达地址10.1.2.2经过的路径。

```
[S1]tracert 10.1.2.2
 traceroute to  10.1.2.2(10.1.2.2), max hops: 30 ,packet length: 40
 1 10.0.113.3 10 ms  1 ms  1 ms
 2 10.0.23.2 40 ms  20 ms  20 ms
```

此时S1选择使用从R3上学习到的默认路由，即通过主用线路访问10.1.2.2。

关闭R3的S2/0/0，模拟公司到运营商的线路故障。

```
[R3]interface s2/0/0
[R3-Serial2/0/0]shutdown
```

待路由收敛后在S1上查看路由表。并检查到10.1.2.2的连通性。

```
[S1]display ip routing-table
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Routing Tables: Public
       Destinations : 13      Routes : 13

Destination/Mask    Proto  Pre  Cost      Flags NextHop         Interface

        0.0.0.0/0   O_ASE  150  21          D   10.0.113.3      Vlanif13
      10.0.1.11/32  Direct 0    0           D   127.0.0.1       InLoopBack0
       10.0.3.3/32  OSPF   10   1           D   10.0.113.3      Vlanif13
       10.0.4.4/32  OSPF   10   1           D   10.0.114.4      Vlanif14
       10.0.5.5/32  OSPF   10   1           D   10.0.115.5      Vlanif15
```

```
   10.0.113.0/24  Direct 0    0           D   10.0.113.1      Vlanif13
   10.0.113.1/32  Direct 0    0           D   127.0.0.1       InLoopBack0
   10.0.114.0/24  Direct 0    0           D   10.0.114.1      Vlanif14
   10.0.114.1/32  Direct 0    0           D   127.0.0.1       InLoopBack0
   10.0.115.0/24  Direct 0    0           D   10.0.115.1      Vlanif15
   10.0.115.1/32  Direct 0    0           D   127.0.0.1       InLoopBack0
    127.0.0.0/8   Direct 0    0           D   127.0.0.1       InLoopBack0
    127.0.0.1/32  Direct 0    0           D   127.0.0.1       InLoopBack0

[S1]ping 10.1.2.2
 PING 10.1.2.2: 56  data bytes, press CTRL_C to break
   Request time out
   Request time out
   Request time out
   Request time out
   Request time out

 --- 10.1.2.2 ping statistics ---
   5 packet(s) transmitted
   0 packet(s) received
   100.00% packet loss
```

可以观察到S1的路由表没有变化，仍旧是通过R3访问目标网络。

由于是上联链路故障，而下联链路中S1是通过比较R3与R5下发默认路由的cost值的大小，最终选择R3下发的默认路由，二者互不影响，所以网络无法正常工作。

恢复R3的S2/0/0，关闭R3的G0/0/1，模拟R3的下联线路故障。

查看路由收敛情况，并检查连通性。

```
[R3]interface s2/0/0
[R3-Serial2/0/0]undo shutdown

[R3]interface g0/0/1
[R3-GigabitEthernet0/0/1]shutdown

[S1]display ip routing-table
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Routing Tables: Public
        Destinations : 10       Routes : 10

Destination/Mask    Proto  Pre  Cost      Flags NextHop         Interface
```

```
     0.0.0.0/01    O_ASE  150  41           D   10.0.115.5      Vlanif15
    10.0.1.11/32   Direct 0    0            D   127.0.0.1       InLoopBack0
     10.0.4.4/32   OSPF   10   1            D   10.0.114.4      Vlanif14
     10.0.5.5/32   OSPF   10   1            D   10.0.115.5      Vlanif15
   10.0.114.0/24   Direct 0    0            D   10.0.114.1      Vlanif14
   10.0.114.1/32   Direct 0    0            D   127.0.0.1       InLoopBack0
   10.0.115.0/24   Direct 0    0            D   10.0.115.1      Vlanif15
   10.0.115.1/32   Direct 0    0            D   127.0.0.1       InLoopBack0
    127.0.0.0/8    Direct 0    0            D   127.0.0.1       InLoopBack0
    127.0.0.1/32   Direct 0    0            D   127.0.0.1       InLoopBack0

[S1]ping 10.1.2.2
  PING 10.1.2.2: 56  data bytes, press CTRL_C to break
    Reply from 10.1.2.2: bytes=56 Sequence=1 ttl=254 time=1 ms
    Reply from 10.1.2.2: bytes=56 Sequence=2 ttl=254 time=1 ms
    Reply from 10.1.2.2: bytes=56 Sequence=3 ttl=254 time=1 ms
    Reply from 10.1.2.2: bytes=56 Sequence=4 ttl=254 time=1 ms
    Reply from 10.1.2.2: bytes=56 Sequence=5 ttl=254 time=1 ms

  --- 10.1.2.2 ping statistics ---
    5 packet(s) transmitted
    5 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 1/1/1 ms
```

此时S1通过R5学习到默认路由，即通过备用链路访问目标网络。

恢复R3的G0/0/1端口。

```
[R3]interface g0/0/1
[R3-GigabitEthernet0/0/1]undo shutdown
```

步骤四. 配置使用缺省路由和部分过滤路由归属到单个 ISP

配置R3、R4、R5到S1的IBGP邻居关系，并添加**next-hop-local**参数，保证S1能学习到从ISP发来的路由更新消息。

```
[R3]bgp 64512
[R3-bgp]peer 10.0.113.1 as-number 64512
[R3-bgp]peer 10.0.113.1 next-hop-local

[R4]bgp 64512
```

```
[R4-bgp]peer 10.0.114.1 as-number 64512
[R4-bgp]peer 10.0.114.1 next-hop-local

[R5]bgp 64512
[R5-bgp]peer 10.0.115.1 as-number 64512
[R5-bgp]peer 10.0.115.1 next-hop-local

[S1]bgp 64512
[S1-bgp]peer 10.0.113.3 as-number 64512
[S1-bgp]peer 10.0.114.4 as-number 64512
[S1-bgp]peer 10.0.115.5 as-number 64512
```

观察S1是否学习到了10.1.2.0/24和10.2.2.0/24。

```
[S1]display ip routing-table
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Routing Tables: Public
       Destinations : 15      Routes : 15

Destination/Mask    Proto  Pre  Cost      Flags NextHop        Interface

       0.0.0.0/0    O_ASE  150  21          D   10.0.113.3     Vlanif13
     10.0.1.11/32   Direct 0    0           D   127.0.0.1      InLoopBack0
      10.0.3.3/32   OSPF   10   1           D   10.0.113.3     Vlanif13
      10.0.4.4/32   OSPF   10   1           D   10.0.114.4     Vlanif14
      10.0.5.5/32   OSPF   10   1           D   10.0.115.5     Vlanif15
    10.0.113.0/24   Direct 0    0           D   10.0.113.1     Vlanif13
    10.0.113.1/32   Direct 0    0           D   127.0.0.1      InLoopBack0
    10.0.114.0/24   Direct 0    0           D   10.0.114.1     Vlanif14
    10.0.114.1/32   Direct 0    0           D   127.0.0.1      InLoopBack0
    10.0.115.0/24   Direct 0    0           D   10.0.115.1     Vlanif15
    10.0.115.1/32   Direct 0    0           D   127.0.0.1      InLoopBack0
      10.1.2.0/24   BGP    255  0          RD   10.0.113.3     Vlanif13
      10.2.2.0/24   BGP    255  0          RD   10.0.113.3     Vlanif13
     127.0.0.0/8    Direct 0    0           D   127.0.0.1      InLoopBack0
     127.0.0.1/32   Direct 0    0           D   127.0.0.1      InLoopBack0
```

这时我们希望通过BGP来影响选路，在R3上添加路由策略**policy_r3**，过滤掉10.1.2.0/24。

```
[R3]acl number 2001
[R3-acl-basic-2001]rule 0 permit source 10.1.2.0 0.0.0.255
[R3-acl-basic-2001]route-policy policy_r3 deny node 10
```

```
[R3-route-policy]if-match acl 2001
[R3-route-policy]route-policy policy_r3 permit node 20
[R3-route-policy]bgp 64512
[R3-bgp]peer 10.0.113.1 route-policy policy_r3 export
```

在R5上添加路由策略**policy_r5**，过滤掉10.2.2.0/24，

```
[R5]acl number 2001
[R5-acl-basic-2001]rule 0 permit source 10.2.2.0 0.0.0.255
[R5-acl-basic-2001]route-policy policy_r5 deny node 10
[R5-route-policy]if-match acl 2001
[R5-route-policy]route-policy policy_r5 permit node 20
[R5-route-policy]bgp 64512
[R5-bgp]peer 10.0.115.1 route-policy policy_r5 export
```

在S1上观察路由表的变化。

```
[S1]display ip routing-table
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Routing Tables: Public
       Destinations : 15      Routes : 15

Destination/Mask    Proto  Pre  Cost      Flags NextHop        Interface

      0.0.0.0/0     O_ASE  150  21          D   10.0.113.3     Vlanif13
    10.0.1.11/32    Direct 0    0           D   127.0.0.1      InLoopBack0
     10.0.3.3/32    OSPF   10   1           D   10.0.113.3     Vlanif13
     10.0.4.4/32    OSPF   10   1           D   10.0.114.4     Vlanif14
     10.0.5.5/32    OSPF   10   1           D   10.0.115.5     Vlanif15
   10.0.113.0/24    Direct 0    0           D   10.0.113.1     Vlanif13
   10.0.113.1/32    Direct 0    0           D   127.0.0.1      InLoopBack0
   10.0.114.0/24    Direct 0    0           D   10.0.114.1     Vlanif14
   10.0.114.1/32    Direct 0    0           D   127.0.0.1      InLoopBack0
   10.0.115.0/24    Direct 0    0           D   10.0.115.1     Vlanif15
   10.0.115.1/32    Direct 0    0           D   127.0.0.1      InLoopBack0
     10.1.2.0/24    BGP    255  0          RD   10.0.115.5     Vlanif15
     10.2.2.0/24    BGP    255  0          RD   10.0.113.3     Vlanif13
    127.0.0.0/8     Direct 0    0           D   127.0.0.1      InLoopBack0
    127.0.0.1/32    Direct 0    0           D   127.0.0.1      InLoopBack0
```

此时到达网络10.1.2.0/24的下一跳是R5，到达网络10.2.2.0/24的下一跳是R3。

关闭R3的S2/0/0。

```
[R3]interface s2/0/0
[R3-Serial2/0/0]shutdown
```

观察S1的路由变化，测试到10.1.2.2的连通性。

```
[S1]display ip routing-table
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Routing Tables: Public
       Destinations : 14      Routes : 14

Destination/Mask    Proto  Pre  Cost      Flags NextHop         Interface

       0.0.0.0/0    O_ASE  150  21          D   10.0.113.3      Vlanif13
     10.0.1.11/32   Direct 0    0           D   127.0.0.1       InLoopBack0
      10.0.3.3/32   OSPF   10   1           D   10.0.113.3      Vlanif13
      10.0.4.4/32   OSPF   10   1           D   10.0.114.4      Vlanif14
      10.0.5.5/32   OSPF   10   1           D   10.0.115.5      Vlanif15
    10.0.113.0/24   Direct 0    0           D   10.0.113.1      Vlanif13
    10.0.113.1/32   Direct 0    0           D   127.0.0.1       InLoopBack0
    10.0.114.0/24   Direct 0    0           D   10.0.114.1      Vlanif14
    10.0.114.1/32   Direct 0    0           D   127.0.0.1       InLoopBack0
    10.0.115.0/24   Direct 0    0           D   10.0.115.1      Vlanif15
    10.0.115.1/32   Direct 0    0           D   127.0.0.1       InLoopBack0
      10.1.2.0/24   BGP    255  0          RD   10.0.115.5      Vlanif15
     127.0.0.0/8    Direct 0    0           D   127.0.0.1       InLoopBack0
     127.0.0.1/32   Direct 0    0           D   127.0.0.1       InLoopBack0
```

此时只有路由10.1.2.0/24，因为R5上做了策略过滤掉了路由10.2.2.0/24。

```
[S1]ping 10.1.2.2
  PING 10.1.2.2: 56  data bytes, press CTRL_C to break
    Reply from 10.1.2.2: bytes=56 Sequence=1 ttl=254 time=1 ms
    Reply from 10.1.2.2: bytes=56 Sequence=2 ttl=254 time=1 ms
    Reply from 10.1.2.2: bytes=56 Sequence=3 ttl=254 time=1 ms
    Reply from 10.1.2.2: bytes=56 Sequence=4 ttl=254 time=1 ms
    Reply from 10.1.2.2: bytes=56 Sequence=5 ttl=254 time=1 ms

  --- 10.1.2.2 ping statistics ---
    5 packet(s) transmitted
    5 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 1/1/1 ms
```

恢复R3的S2/0/0接口。

```
[R3]interface s2/0/0
[R3-Serial2/0/0]undo shutdown
```

步骤五. 配置使用 BGP 路由归属到多个 ISP

这时公司又申请了一条Internet线路接入ISP2，为了采用BGP来进行路径选择，这里先删除前面试验中OSPF发布的默认路由。

```
[R3]ospf
[R3-ospf-1]undo default-route-advertise

[R5]ospf
[R5-ospf-1]undo default-route-advertise
```

删除R3和R5上对路由进行过滤的策略。

```
[R3]undo route-policy policy1

[R5]undo route-policy policy2
```

删除R3和R5上将OSPF路由引入到BGP中的命令。

```
[R3]bgp 64512
[R3-bgp]undo import-route ospf 1

[R5]bgp 64512
[R5-bgp]undo import-route ospf 1
```

在R1和R2、R1和R4之间建立EBGP邻居关系，使ISP2也能传送10.1.2.0/24和10.2.2.0/24。

```
[R1]bgp 200
[R1-bgp]peer 10.0.12.2 as-number 100
[R1-bgp]peer 10.0.14.4 as-number 64512

[R2]bgp 100
[R2-bgp]peer 10.0.12.1 as-number 200

[R4]bgp 64512
[R4-bgp]peer 10.0.14.1 as-number 200
```

在S1上观察路由10.1.2.0/24和10.2.2.0/24，注意当前的选路原则。

```
[S1]display bgp routing-table

 Total Number of Routes: 6

 BGP Local router ID is 10.0.11.11
 Status codes: * - valid, > - best, d - damped,
               h - history,  i - internal, s - suppressed, S - Stale
               Origin : i - IGP, e - EGP, ? - incomplete

     Network            NextHop        MED        LocPrf    PrefVal Path/Ogn

 *>i  10.1.2.0/24    10.0.113.3     0          100        0       100i
 * i                 10.0.115.5     0          100        0       100i
 * i                 10.0.114.4                100        0      200 100i
 *>i  10.2.2.0/24    10.0.113.3     0          100        0       100i
 * i                 10.0.115.5     0          100        0       100i
 * i                 10.0.114.4                100        0      200 100i
```

现在我们希望公司使用连接到ISP2的新线路来访问网络10.2.2.0/24。在R4上设置路由策略**policy_r4**，将该路由的本地优先属性改为150。

```
[R4]acl number 2001
[R4-acl-basic-2001]rule 0 permit source 10.2.2.0 0.0.0.255
[R4-acl-basic-2001]route-policy policy_r4 permit node 10
[R4-route-policy]if-match acl 2001
[R4-route-policy]apply local-preference 150
[R4-route-policy]route-policy policy_r4 permit node 20
```

将该路由策略对S1发布。

```
[R4]bgp 64512
[R4-bgp]peer 10.0.114.1 route-policy policy_r4 export
```

在S1上检查BGP路由表。

```
[S1]display bgp routing-table

 Total Number of Routes: 6

 BGP Local router ID is 10.0.11.11
 Status codes: * - valid, > - best, d - damped,
               h - history,  i - internal, s - suppressed, S - Stale
               Origin : i - IGP, e - EGP, ? - incomplete

     Network            NextHop        MED        LocPrf    PrefVal Path/Ogn
```

```
*>i  10.1.2.0/24    10.0.113.3     0          100        0       100i
* i                 10.0.115.5     0          100        0       100i
* i                 10.0.114.4                100        0       200 100i
*>i  10.2.2.0/24    10.0.114.4                150        0       200 100i
* i                 10.0.113.3     0          100        0       100i
* i                 10.0.115.5     0          100        0       100i
```

此时S1使用通过R4连接的ISP2获得路由到达网段10.2.2.0/24。

关闭R4的S1/0/0端口模拟故障。

```
[R4]interface s1/0/0
[R4-Serial1/0/0]shutdown
```

查看S1上BGP路由表的变化。

```
[S1]display bgp routing-table

 Total Number of Routes: 4

 BGP Local router ID is 10.0.11.11
 Status codes: * - valid, > - best, d - damped,
               h - history,  i - internal, s - suppressed, S - Stale
               Origin : i - IGP, e - EGP, ? - incomplete

     Network            NextHop        MED        LocPrf    PrefVal Path/Ogn

*>i  10.1.2.0/24    10.0.113.3     0          100        0       100i
* i                 10.0.115.5     0          100        0       100i
*>i  10.2.2.0/24    10.0.113.3     0          100        0       100i
* i                 10.0.115.5     0          100        0       100i
```

此时S1通过R3所连接的ISP1获得路由10.1.2.0/24和10.2.2.0/24。

打开R4的S1/0/0接口。

```
[R4]interface s1/0/0
[R4-Serial1/0/0]undo shutdown
```

检查S1上的BGP路由表，查看是否恢复。

```
[S1]display bgp routing-table

 Total Number of Routes: 6
```

```
BGP Local router ID is 10.0.11.11
Status codes: * - valid, > - best, d - damped,
            h - history,  i - internal, s - suppressed, S - Stale
            Origin : i - IGP, e - EGP, ? - incomplete

    Network          NextHop        MED        LocPrf    PrefVal Path/Ogn

*>i  10.1.2.0/24    10.0.113.3     0          100        0      100i
* i                 10.0.115.5     0          100        0      100i
* i                 10.0.114.4                100        0      200 100i
*>i  10.2.2.0/24    10.0.114.4                150        0      200 100i
* i                 10.0.113.3     0          100        0      100i
* i                 10.0.115.5     0          100        0      100i
```

附加实验：思考并验证

思考在步骤三中，关闭R3的S2/0/0接口后，虽然公司到运营商的主用链路发生故障，但是此时R5与ISP1之间的备用链路仍然正常，该如何解决此时的连通问题？

在这个例子中，多归属到两个运营商，怎样实现对同一网段入流量的负载分担？

最终设备配置

```
<R1>display current-configuration
[V200R001C00SPC200]
#
 sysname R1
#
interface Serial1/0/0
 link-protocol ppp
 ip address 10.0.12.1 255.255.255.0
#
interface Serial3/0/0
 link-protocol ppp
 ip address 10.0.14.1 255.255.255.0
#
interface LoopBack0
 ip address 10.0.1.1 255.255.255.255
#
bgp 200
```

```
 peer 10.0.12.2 as-number 100
 peer 10.0.14.4 as-number 64512
 #
 ipv4-family unicast
  undo synchronization
  maximum load-balancing 4
  peer 10.0.12.2 enable
  peer 10.0.14.4 enable
#
return

<R2>display current-configuration
[V200R001C00SPC200]
#
 sysname R2
#
interface Serial1/0/0
 link-protocol ppp
 ip address 10.0.12.2 255.255.255.0
#
interface Serial2/0/0
 link-protocol ppp
 ip address 10.0.23.2 255.255.255.0
#
interface GigabitEthernet0/0/0
 ip address 10.0.25.2 255.255.255.0
#
interface LoopBack0
 ip address 10.0.2.2 255.255.255.255
#
interface LoopBack1
 ip address 10.1.2.2 255.255.255.0
#
interface LoopBack2
 ip address 10.2.2.2 255.255.255.0
#
bgp 100
 peer 10.0.12.1 as-number 200
 peer 10.0.23.3 as-number 64512
 peer 10.0.25.5 as-number 64512
 #
 ipv4-family unicast
  undo synchronization
```

```
  network 10.1.2.0 255.255.255.0
  network 10.2.2.0 255.255.255.0
  maximum load-balancing 4
  peer 10.0.12.1 enable
  peer 10.0.23.3 enable
  peer 10.0.25.5 enable
#
return

<R3>display current-configuration
[V200R001C00SPC200]
#
 sysname R3
#
acl number 2001
 rule 0 permit source 10.1.2.0 0.0.0.255
#
interface Serial2/0/0
 link-protocol ppp
 ip address 10.0.23.3 255.255.255.0
#
interface GigabitEthernet0/0/1
 ip address 10.0.113.3 255.255.255.0
#
interface LoopBack0
 ip address 10.0.3.3 255.255.255.255
#
bgp 64512
 peer 10.0.23.2 as-number 100
 peer 10.0.113.1 as-number 64512
 #
 ipv4-family unicast
  undo synchronization
  maximum load-balancing 4
  peer 10.0.23.2 enable
  peer 10.0.113.1 enable
  peer 10.0.113.1 next-hop-local
#
ospf 1
 area 0.0.0.0
  network 10.0.113.3 0.0.0.0
  network 10.0.3.3 0.0.0.0
#
```

```
return

<R4>display current-configuration
[V200R001C00SPC200]
#
 sysname R4
#
interface Serial1/0/0
 link-protocol ppp
 ip address 10.0.14.4 255.255.255.0
#
interface GigabitEthernet0/0/1
 ip address 10.0.114.4 255.255.255.0
#
interface LoopBack0
 ip address 10.0.4.4 255.255.255.255
#
bgp 64512
 peer 10.0.14.1 as-number 200
 peer 10.0.114.1 as-number 64512
 #
 ipv4-family unicast
  undo synchronization
  maximum load-balancing 4
  peer 10.0.14.1 enable
  peer 10.0.114.1 enable
  peer 10.0.114.1 route-policy policy_r4 export
  peer 10.0.114.1 next-hop-local
#
ospf 1
 area 0.0.0.0
  network 10.0.114.4 0.0.0.0
  network 10.0.4.4 0.0.0.0
#
route-policy policy_r4 permit node 10
 if-match acl 2001
 apply local-preference 150
route-policy policy_r4 permit node 20
#
Return

<R5>display current-configuration
[V200R001C00SPC200]
```

```
#
 sysname R5
#
interface GigabitEthernet0/0/0
 ip address 10.0.25.5 255.255.255.0
#
interface GigabitEthernet0/0/1
 ip address 10.0.115.5 255.255.255.0
#
interface LoopBack0
 ip address 10.0.5.5 255.255.255.255
#
bgp 64512
 peer 10.0.25.2 as-number 100
 peer 10.0.115.1 as-number 64512
 #
 ipv4-family unicast
  undo synchronization
  maximum load-balancing 4
  peer 10.0.25.2 enable
  peer 10.0.115.1 enable
  peer 10.0.115.1 next-hop-local
#
ospf 1
 area 0.0.0.0
  network 10.0.115.5 0.0.0.0
  network 10.0.5.5 0.0.0.0
#
return

<S1>display current-configuration
#
!Software Version V100R005C01SPC100
 sysname S1
#
interface Vlanif13
 ip address 10.0.113.1 255.255.255.0
#
interface Vlanif14
 ip address 10.0.114.1 255.255.255.0
#
interface Vlanif15
 ip address 10.0.115.1 255.255.255.0
```

```
#
interface GigabitEthernet0/0/3
 port link-type access
 port default vlan 13
#
interface GigabitEthernet0/0/4
 port link-type access
 port default vlan 14
#
interface GigabitEthernet0/0/5
 port link-type access
 port default vlan 15
#
interface LoopBack0
 ip address 10.0.1.11 255.255.255.255
#
bgp 64512
 peer 10.0.113.3 as-number 64512
 peer 10.0.114.4 as-number 64512
 peer 10.0.115.5 as-number 64512
 #
 ipv4-family unicast
  undo synchronization
  peer 10.0.113.3 enable
  peer 10.0.114.4 enable
  peer 10.0.115.5 enable
#
ospf 1
 area 0.0.0.0
  network 10.0.113.1 0.0.0.0
  network 10.0.114.1 0.0.0.0
  network 10.0.115.1 0.0.0.0
  network 10.0.1.11 0.0.0.0
#
return
```

实验 3-6 BGP 故障排除

学习目的

- 掌握BGP邻居无法建立的故障排除
- 掌握BGP的硬复位的作用
- 掌握BGP的软复位的作用
- 掌握BGP相关Debug命令的使用

拓扑图

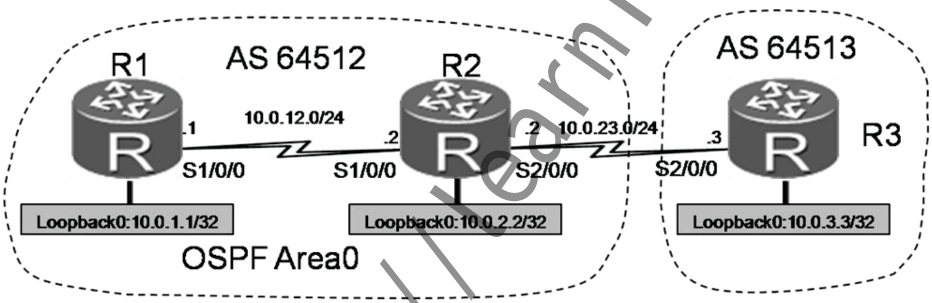

图3-6 BGP故障排除

场景

你是公司的网络管理员。公司的网络采用了BGP协议作为路由协议。公司的网络由多个自治系统组成，不同的分支机构使用了不同的AS号。现在你已经完成公司网络的搭建工作，可是在配置BGP的过程中你遇到了不少问题，最后你通过故障排除的思想和方法，你成功的找到了各种错误，并进行了网络的故障排除。

学习任务

步骤一. 基础配置与 IP 编址

首先给所有路由器配置物理接口及Loopback接口的IP地址和掩码，各Loopback接口均为32位掩码，地址的规划如拓扑图中所示。

```
<R1>system-view
Enter system view, return user view with Ctrl+Z.
[R1]interface Serial 1/0/0
[R1-Serial1/0/0]ip address 10.0.12.1 24
[R1-Serial1/0/0]quit
[R1]interface LoopBack 0
[R1-LoopBack0]ip add 10.0.1.1 32
[R1-LoopBack0]quit

[R2]interface Serial 1/0/0
[R2-Serial1/0/0]ip address 10.0.12.2 24
[R2-Serial1/0/0]quit
[R2]interface Serial 2/0/0
[R2-Serial2/0/0]ip address 10.0.23.2 24
[R2-Serial2/0/0]quit
[R2]interface LoopBack 0
[R2-LoopBack0]ip address 10.0.2.2 32
[R2-LoopBack0]quit

[R3]interface Serial 2/0/0
[R3-Serial2/0/0]ip address 10.0.23.3 24
[R3-Serial2/0/0]quit
[R3]interface LoopBack 0
[R3-LoopBack0]ip address 10.0.3.3 32
[R3-LoopBack0]quit
```

配置完成后，测试直连链路的连通性。

```
[R2]ping -c 1 10.0.12.1
  PING 10.0.12.1: 56  data bytes, press CTRL_C to break
    Reply from 10.0.12.1: bytes=56 Sequence=1 ttl=255 time=40 ms

  --- 10.0.12.1 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 40/40/40 ms

[R2]ping -c 1 10.0.23.3
  PING 10.0.23.3: 56  data bytes, press CTRL_C to break
    Reply from 10.0.23.3: bytes=56 Sequence=1 ttl=255 time=38 ms

  --- 10.0.23.3 ping statistics ---
    1 packet(s) transmitted
```

```
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 38/38/38 ms
```

结果显示直连链路连通性正常。

步骤二. 配置 IGP 及 BGP

在这个场景中，在AS 64512内部我们使用OSPF作为IGP，所有设备属于区域0。各路由器使用Loopback0的地址作为Router ID。R1的S1/0/0和Loopback 0所连接的网段运行OSPF。

```
[R1]ospf 1 router-id 10.0.1.1
[R1-ospf-1]area 0
[R1-ospf-1-area-0.0.0.0]network 10.0.12.1 0.0.0.0
[R1-ospf-1-area-0.0.0.0]network 10.0.1.1 0.0.0.0
```

R2的S1/0/0和Loopback 0所连接的网段运行OSPF。

```
[R2]ospf 1 router-id 10.0.2.2
[R2-ospf-1]area 0
[R2-ospf-1-area-0.0.0.0]network 10.0.12.2 0.0.0.0
[R2-ospf-1-area-0.0.0.0]network 10.0.2.2 0.0.0.0
```

配置完成以后检查R1到R2的Loopback0地址的连通性。

```
[R1]ping -c 1 -a 10.0.1.1 10.0.2.2
  PING 10.0.2.2: 56  data bytes, press CTRL_C to break
    Reply from 10.0.2.2: bytes=56 Sequence=1 ttl=255 time=40 ms

  --- 10.0.2.2 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 40/40/40 ms
```

在R1、R2之间配置IBGP，在R2和R3之间配置EBGP，均使用Loopback地址建立对等体关系，为保证路由信息的正常传递，在R2上针对R1配置**Next-hop-local**，在R3上你错误的将对等体10.0.2.2的AS号配置为64514了。

```
[R1]bgp 64512
[R1-bgp]peer 10.0.2.2 as-number 64512

[R2]bgp 64512
```

```
[R2-bgp]peer 10.0.1.1 as-number 64512
[R2-bgp]peer 10.0.1.1 next-hop-local
[R2-bgp]peer 10.0.3.3 as-number 64513

[R3]bgp 64513
[R3-bgp]peer 10.0.2.2 as-number 64514
```

步骤三. 排除对等建立的故障

在完成上面的配置以后，你发现路由器之间对等体关系并没有建立起来，首先在R2上观察对等体关系。

```
[R2]display bgp peer

 BGP local router ID : 10.0.12.2
 Local AS number : 64512
 Total number of peers : 2              Peers in established state : 0

  Peer            V          AS  MsgRcvd  MsgSent  OutQ  Up/Down       State  PrefRcv

  10.0.1.1        4       64512        0        0     0  00:05:36     Active        0
  10.0.3.3        4       64513        0        0     0  00:05:21       Idle        0
```

从上面的输出中我们可以看到，10.0.1.1的状态为Active，10.0.2.2的状态为Idle。BGP对等体关系建立正常的情况下，State的状态应为Established。长时间停留在其他状态均为故障的现象，下面将逐步排除各种故障。

一般情况下，当对等体IP地址对于本地路由器是不可达的时，对等体状态会显示为Idle。也就是路由器根本没有向对等体发起TCP连接。当对等体IP地址可达，但TCP连接建立存在问题时，就可以看到对等体之间停留在Active状态。

首先考虑R2与R3之间的对等体关系的问题，先检查R2和R3之间Loopback地址之间的连通性。

```
[R2]ping -c 1 -a 10.0.2.2 10.0.3.3
  PING 10.0.3.3: 56  data bytes, press CTRL_C to break
    Request time out

  --- 10.0.3.3 ping statistics ---
    1 packet(s) transmitted
    0 packet(s) received
    100.00% packet loss
```

发现R2与R3的Loopback0接口地址之间的通讯存在问题。

检查R2的路由表。

```
[R2]display ip routing-table
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Routing Tables: Public
        Destinations : 14       Routes : 14

Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface

      10.0.1.1/32   OSPF    10   1562        D   10.0.12.1       Serial1/0/0
      10.0.2.2/32   Direct  0    0           D   127.0.0.1       InLoopBack0
     10.0.12.0/24   Direct  0    0           D   10.0.12.2       Serial1/0/0
     10.0.12.1/32   Direct  0    0           D   10.0.12.1       Serial1/0/0
     10.0.12.2/32   Direct  0    0           D   127.0.0.1       InLoopBack0
   10.0.12.255/32   Direct  0    0           D   127.0.0.1       InLoopBack0
     10.0.23.0/24   Direct  0    0           D   10.0.23.2       Serial2/0/0
     10.0.23.2/32   Direct  0    0           D   127.0.0.1       InLoopBack0
     10.0.23.3/32   Direct  0    0           D   10.0.23.3       Serial2/0/0
   10.0.23.255/32   Direct  0    0           D   127.0.0.1       InLoopBack0
     127.0.0.0/8    Direct  0    0           D   127.0.0.1       InLoopBack0
     127.0.0.1/32   Direct  0    0           D   127.0.0.1       InLoopBack0
127.255.255.255/32  Direct  0    0           D   127.0.0.1       InLoopBack0
255.255.255.255/32  Direct  0    0           D   127.0.0.1       InLoopBack0
```

发现没有到达地址10.0.3.3的路由。

检查R3的路由表。

```
[R3]display ip routing-table
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Routing Tables: Public
        Destinations : 9        Routes : 9

Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface

      10.0.3.3/32   Direct  0    0           D   127.0.0.1       InLoopBack0
     10.0.23.0/24   Direct  0    0           D   10.0.23.3       Serial2/0/0
     10.0.23.2/32   Direct  0    0           D   10.0.23.2       Serial2/0/0
     10.0.23.3/32   Direct  0    0           D   127.0.0.1       InLoopBack0
   10.0.23.255/32   Direct  0    0           D   127.0.0.1       InLoopBack0
     127.0.0.0/8    Direct  0    0           D   127.0.0.1       InLoopBack0
```

```
     127.0.0.1/32   Direct  0    0          D   127.0.0.1       InLoopBack0
127.255.255.255/32  Direct  0    0          D   127.0.0.1       InLoopBack0
255.255.255.255/32  Direct  0    0          D   127.0.0.1       InLoopBack0
```

也没有到达R2的Loopback0地址10.0.2.2的路由。

对于不同的AS，使用静态路由可以实现相邻路由器的Loopback接口地址之间的可达性。

在这里，我们需要在R2和R3上分别添加到对端Loopback接口连接网段的静态路由。

```
[R2]ip route-static 10.0.3.3 32 10.0.23.3

[R3]ip route-static 10.0.2.2 32 10.0.23.2
```

检查R2到R3之间的连通性。

```
[R2]ping -c 1 -a 10.0.2.2 10.0.3.3
  PING 10.0.3.3: 56  data bytes, press CTRL_C to break
    Reply from 10.0.3.3: bytes=56 Sequence=1 ttl=255 time=30 ms

  --- 10.0.3.3 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 30/30/30 ms
```

检查R2的BGP对等体关系。

```
[R2]display bgp peer

 BGP local router ID : 10.0.12.2
 Local AS number : 64512
 Total number of peers : 2              Peers in established state : 0

  Peer            V        AS  MsgRcvd  MsgSent  OutQ  Up/Down      State PrefRcv

  10.0.1.1        4     64512        0        0     0 05:23:27     Active       0
  10.0.3.3        4     64513        0        0     0 05:23:02     Active       0
```

R2和R3之间的状态由刚才的Idle状态转变为Active状态。

接下来考虑R1和R2之间的对等体关系，刚才在配置OSPF之后已经验证了R1的Loopback接口地址到R2的Loopback接口地址之间的连通性。

BGP是通过TCP 179端口进行通信的。首先检查路由器的179端口是否处于打开状态。

在R1和R2上分别检查TCP状态。

```
[R1]display tcp status
TCPCB     Tid/Soid  Local Add:port       Foreign Add:port      VPNID  State
194b9500  8  /2      0.0.0.0:22            0.0.0.0:0           23553   Listening
194b939c  8  /1      0.0.0.0:23            0.0.0.0:0           23553   Listening
194b90d4  106/1      0.0.0.0:80            0.0.0.0:0           0       Listening
194b9a90  234/2      0.0.0.0:179           10.0.2.2:0          0       Listening
194b9664  8  /3      0.0.0.0:830           0.0.0.0:0           23553   Listening
194b9238  6  /1      0.0.0.0:7547          0.0.0.0:0           0       Listening

[R2]display tcp status
TCPCB     Tid/Soid  Local Add:port        Foreign Add:port     VPNID  State
1949a048  234/5      0.0.0.0:0             0.0.0.0:0           0       Closed
19499d80  8  /2      0.0.0.0:22            0.0.0.0:0           23553   Listening
19499c1c  8  /1      0.0.0.0:23            0.0.0.0:0           23553   Listening
19499954  106/1      0.0.0.0:80            0.0.0.0:0           0       Listening
1949a474  234/2      0.0.0.0:179           10.0.1.1:0          0       Listening
1949a310  234/4      0.0.0.0:179           10.0.3.3:0          0       Listening
19499ee4  8  /3      0.0.0.0:830           0.0.0.0:0           23553   Listening
19499ab8  6  /1      0.0.0.0:7547          0.0.0.0:0           0       Listening
```

可以看到，对等体之间对应地址的179端口均处于监听状态。显示BGP协议在单个路由器上工作正常。

然后通过**debug**命令在R1上观察是否收到R2发来的BGP数据包。

```
<R1>terminal debugging
<R1>debugging tcp packet
Dec  7 2011 10:08:16.620.1+00:00 R1 SOCKET/7/TCP PACKET:
TCP debug packet information:
1323252496: Input: no port,
(src = 10.0.12.2:52688,dst = 10.0.1.1:179,VrfIndex = 0,seq = 2254758724,
ack = 0,datalen = 0,optlen = 4,flag =  SYN ,window = 16384,ttl = 0,tos = 0,MSS
= 0)

Dec  7 2011 10:08:16.620.2+00:00 R1 SOCKET/7/TCP PACKET:
TCP debug packet information:
1323252496: Output: task = (0), socketid = 0,
(src = 10.0.1.1:179,dst = 10.0.12.2:52688,VrfIndex = 0,seq = 0,
ack = 2254758725,datalen = 0,optlen = 0,flag =  ACK  RST ,window = 0,ttl = 255,tos
= 0,MSS = 0)
```

从上面的输出中可以看到R2发来的目的端口号为179的数据包的源地址为10.0.12.2。查看拓扑后发现该地址为R2的Serial1/0/0的接口地址。

在建立对等体关系时，使用的是R2的Loopback接口地址，这就造成R1和R2之间的对等体关系建立不起来。所以我们需要在建立对等体关系时需使用**connect-interface**指定更新源地址。

同样的，R2和R3之间也存在这个问题。在这里，一并进行修改。

```
[R1]bgp 64512
[R1-bgp]peer 10.0.2.2 connect-interface LoopBack 0

[R2]bgp 64512
[R2-bgp]peer 10.0.1.1 connect-interface LoopBack 0
[R2-bgp]peer 10.0.3.3 connect-interface LoopBack 0

[R3]bgp 64513
[R3-bgp]peer 10.0.2.2 connect-interface LoopBack 0
```

修改完成以后再次在R2上检查对等体关系。

```
[R2]display bgp peer

 BGP local router ID : 10.0.12.2
 Local AS number : 64512
 Total number of peers : 2          Peers in established state : 1

  Peer            V          AS  MsgRcvd  MsgSent  OutQ  Up/Down       State  PrefRcv

  10.0.1.1        4       64512       16       17     0 00:14:18 Established       0
  10.0.3.3        4       64513        0        0     0 00:14:35      Active       0
```

看到R1与R2之间的状态已经是Established状态了。

在R3上使用**debug**命令测试是否收到BGP数据包，并检查数据包的内容。

```
<R3>terminal debugging
<R3>debugging ip packet
Dec  7 2011 10:51:44.30.5+00:00 R3 IP/7/debug_case:
Delivering, interface = S2/0/0, version = 4, headlen = 20, tos = 192,
pktlen = 40, pktid = 4752, offset = 0, ttl = 1, protocol = 6,
checksum = 36220, s = 10.0.2.2, d = 10.0.3.3
prompt: Packet is before IP_Reass before really deliver to up.

Dec  7 2011 10:51:44.30.6+00:00 R3 IP/7/debug_case:
```

```
Sending, interface = S2/0/0, version = 4, headlen = 20, tos = 0,
pktlen = 40, pktid = 9953, offset = 0, ttl = 255, protocol = 6,
checksum = 31722, s = 10.0.3.3, d = 10.0.2.2
prompt: Sending the packet from local at S2/0/0
```

注意到R3收到的BGP数据包TTL值为1。对于EBGP来说，路由器向对等体发送的数据包默认的TTL就是1。

在这个场景中，R2和R3之间使用Loopback地址建立对等体关系。从R2的Loopback地址到R3的Loopback地址有2跳。所以该数据包在还没有到达R2的Loopback地址之前就因为TTL超时被丢弃了。

为了解决这个问题，我们需要修改EBGP对等体之间发送的数据包的TTL值。

```
[R2]bgp 64512
[R2-bgp]peer 10.0.3.3 ebgp-max-hop 2

[R3]bgp 64513
[R3-bgp]peer 10.0.2.2 ebgp-max-hop 2
```

配置完成以后，在R2上重新检查对等体之间的关系。

```
[R2]display bgp peer

 BGP local router ID : 10.0.12.2
 Local AS number : 64512
 Total number of peers : 2           Peers in established state : 1

  Peer            V          AS  MsgRcvd  MsgSent  OutQ  Up/Down       State
PrefRcv

  10.0.1.1        4       64512        3        4     0 00:01:34 Established       0
  10.0.3.3        4       64513        0        1     0 00:00:44      Active       0
```

R2和R3之间仍停留在Active状态。

在R3上检查BGP的错误。

```
[R3]display bgp error
 Error Type     : Peer Error
 Date/Time      : 2011/12/07 11:24:37
 Peer Address   : 10.0.2.2
 VRF Name       : Public
 Error Info     : Incorrect remote AS

 Error Type     : Peer Error
```

```
 Date/Time     : 2011/12/07 11:25:09
 Peer Address  : 10.0.2.2
 VRF Name      : Public
 Error Info    : Incorrect remote AS

 Error Type    : Peer Error
 Date/Time     : 2011/12/07 11:25:41
 Peer Address  : 10.0.2.2
 VRF Name      : Public
 Error Info    : Incorrect remote AS

<R3>terminal debugging
<R3>debugging bgp packet verbose
Dec  7 2011 11:31:01.540.1+00:00 R3 RM/6/RMDEBUG:
 BGP.Public: Err/SubErr: 2/2 Errdata: 41040000fc00
 Identified in OPEN MSG from 10.0.2.2.

Dec  7 2011 11:31:01.540.2+00:00 R3 RM/6/RMDEBUG:

Dec  7 2011 11:31:01.540.3+00:00 R3 RM/6/RMDEBUG:
 BGP.Public: Err/SubErr: 2/2 Errdata: 41040000fc00
 Identified in OPEN MSG from 10.0.2.2.
```

出现的是AS号错误的消息。

使用**debug**命令来诊断该错误。

从上面的输出中我们可以看到，错误号/子错误号是2。翻阅文档可知，该错误就是AS号不匹配。接下来我们在R3上修改对等体的AS号。

```
[R3]bgp 64513
[R3-bgp]undo peer 10.0.2.2
[R3-bgp]peer 10.0.2.2 as-number 64512
[R3-bgp]peer 10.0.2.2 ebgp-max-hop 2
[R3-bgp]peer 10.0.2.2 connect-interface LoopBack0
```

这时再检查R2和R3之间的对等体关系。

```
[R2]display bgp peer

 BGP local router ID : 10.0.12.2
 Local AS number : 64512
 Total number of peers : 2          Peers in established state : 2

  Peer         V       AS  MsgRcvd  MsgSent  OutQ  Up/Down       State PrefRcv
```

```
 10.0.1.1      4      64512       81       82      0 01:19:18 Established      0
 10.0.3.3      4      64513        3        4      0 00:01:12 Established      0
```

步骤四. BGP 安全

BGP通常用在骨干网上，其安全性显得尤为重要。一旦BGP路由器被攻击，将会造成大面积网络瘫痪。

通常为了防止非法恶意用户冒充合法路由器与BGP路由器建立对等体关系，在BGP对等体会话之间会设置MD5认证。

R1和R2之间开启MD5的认证。首先设置一个错误的密码。在R1上设置密码为huawei，在R2上设置密码为123，观察对等体关系的变化情况。

```
[R1]bgp 64512
[R1-bgp]peer 10.0.2.2 password simple huawei

[R2]bgp 64512
[R2-bgp]peer 10.0.1.1 password simple 123
```

在R1上重置BGP对等体关系，这时可以观察到R1与R2的对等体关系已停留在Connect和Active状态。无法进入Established状态。

```
<R1>reset bgp 10.0.2.2
[R1]display bgp peer

 BGP local router ID : 10.0.12.1
 Local AS number : 64512
 Total number of peers : 1             Peers in established state : 0

  Peer            V          AS  MsgRcvd  MsgSent  OutQ  Up/Down       State PrefRcv

  10.0.2.2        4       64512        0        0     0 00:03:39     Connect       0
```

我们把R2的密码改成huawei。

```
[R2-bgp]undo peer 10.0.1.1 password
[R2-bgp]peer 10.0.1.1 password simple huawei
```

等待约半分钟，再次查看对等体关系。

```
[R2]display bgp peer
```

```
 BGP local router ID : 10.0.12.2
 Local AS number : 64512
 Total number of peers : 2              Peers in established state : 2

  Peer            V          AS  MsgRcvd  MsgSent  OutQ  Up/Down       State  PrefRcv

  10.0.1.1        4       64512        2        2     0 00:00:34 Established       0
  10.0.3.3        4       64513      166      167     0 02:44:05 Established       0
```

这时，R1和R2之间的对等体关系已经到达了Established状态。

在这个场景中AS 64512的管理员不希望AS64513中的路由器看自己的真实AS号。

fake-as的功能就能达到这个目的，该命令可以用来为本端对等体指定一个伪AS号。

我们在R2上针对R3配置这条命令，伪装以后的AS号为100。

同时，也要修改R3配置的BGP对等体R2的AS号。

```
[R2]bgp 64512
[R2-bgp]peer 10.0.3.3 fake-as 100

[R3]bgp 64513
[R3-bgp]undo peer 10.0.2.2
[R3-bgp]peer 10.0.2.2 as-number 100
[R3-bgp]peer 10.0.2.2 ebgp-max-hop 2
[R3-bgp]peer 10.0.2.2 connect-interface LoopBack0
```

在R3上检查对等体，发现R2的AS号已经变成了100。

```
[R3]display bgp peer

 BGP local router ID : 10.0.23.3
 Local AS number : 64513
 Total number of peers : 1              Peers in established state : 1

  Peer            V         AS  MsgRcvd  MsgSent  OutQ  Up/Down        State   PrefRcv

  10.0.2.2        4         100        2        2     0 00:00:28   Established      0
```

在R2上发布自己的Loopback 0接口所在网段，观察在R3上学习到的BGP路由的AS-Path属性。

```
[R2]bgp 64512
[R2-bgp]network 10.0.2.2 32
```

```
[R3]display bgp routing-table

 BGP Local router ID is 10.0.23.3
 Status codes: * - valid, > - best, d - damped,
               h - history,  i - internal, s - suppressed, S - Stale
               Origin : i - IGP, e - EGP, ? - incomplete


 Total Number of Routes: 1
      Network            NextHop        MED        LocPrf    PrefVal Path/Ogn

    10.0.2.2/32        10.0.2.2        0                      0      100i
```

可以看到在R3上学习到的路由10.0.2.2/32的AS-Path是100，即R3认为该路由是在AS100始发的，这样就达到了隐藏AS 64512的目的。

BGP还提供了一种安全特性：GTSM。该特性通过检测IP报文头中的TTL值是否在一个预先定义好的范围内来对路由器进行保护。也就是说，若收到的BGP报文的TTL值超出了预先设定的值的范围以后，就丢弃该报文。注意GTSM和ebgp-max-hop功能均会影响到发送出去的BGP报文的TTL值，存在冲突。只能对同一对等体或对等体组使能两种功能中的一种。

在这个场景中，我们对R2和R3之间的链路开启该特性，观察BGP报文的交互。首先在R2的系统视图下设置违背GTSM规则的缺省动作，在这里我们选择**drop**，即不符合要求的BGP数据包将会被丢弃。

```
[R2]gtsm default-action drop
```

然后在BGP视图下针对R3开启GTSM特性。在配置之前需先要删除**ebgp-max-hop**的配置，注意这里从R2和R3是直连的，所以这里的**valid-ttl-hops**参数填1。

```
[R2]bgp 64512
[R2-bgp]undo peer 10.0.3.3 ebgp-max-hop
[R2-bgp]peer 10.0.3.3 valid-ttl-hops 1
[R2-bgp]peer 10.0.1.1 valid-ttl-hops 1
```

在R1和R3上进行相同的操作。

```
[R1]gtsm default-action drop
[R1]bgp 64512
[R1-bgp]peer 10.0.2.2 valid-ttl-hops 1

[R3]gtsm default-action drop
```

```
[R3]bgp 64513
[R3-bgp]undo  peer 10.0.2.2 ebgp-max-hop
[R3-bgp]peer 10.0.2.2 valid-ttl-hops 1
```

检查R2和R3之间的对等体关系。

```
[R3]dis bgp peer

 BGP local router ID : 10.0.23.3
 Local AS number : 64513
 Total number of peers : 1              Peers in established state : 1

  Peer            V          AS  MsgRcvd  MsgSent  OutQ  Up/Down       State PrefRcv

  10.0.2.2        4         100        3        2     0 00:00:06 Established     1
```

在R3上观察使用了该命令以后BGP报文TTL值的变化。

```
<R3>debugging ip packet
Dec  7 2011 16:34:51.10.1+00:00 R3 IP/7/debug_case:
Receiving, interface = S2/0/0, version = 4, headlen = 20, tos = 192,
pktlen = 59, pktid = 8820, offset = 0, ttl = 255, protocol = 6,
checksum = 32644, s = 10.0.2.2, d = 10.0.3.3
prompt: Receiving IP packet from S2/0/0

Dec  7 2011 16:34:51.10.2+00:00 R3 IP/7/debug_case:
Receiving, interface = Serial2/0/0, version = 4, headlen = 20, tos = 192,
pktlen = 59, pktid = 8820, offset = 0, ttl = 255, protocol = 6,
checksum = 32644, s = 10.0.2.2, d = 10.0.3.3
prompt: IP_ProcessByBoard Begin!
```

这时看到R3接收到的从R2发来的数据包的TTL值为255，而不是EBGP的默认值1。为了确认GTSM会丢弃TTL不合规定的BGP数据包，我们首先在R3上打开GTSM日志记录功能，当有数据包被GTSM丢弃时就记录相应信息。

```
[R3]gtsm log drop-packet all
```

然后在R2上配置**ebgp-max-hop**命令，使R2发给R3的BGP数据包的TTL值小于254。

```
[R2]bgp 64512
[R2-bgp]undo peer 10.0.3.3 valid-ttl-hops
[R2-bgp]peer 10.0.3.3 ebgp-max-hop 253
```

等待一段时间以后，可以看到R2和R3之间的对等体关系进入了IDLE状态。

现在R3上查看GTSM统计信息可以发现，已经有数据包被GTSM丢弃了。

```
Dec  7 2011 16:48:34+00:00 R3 %%01BGP/3/STATE_CHG_UPDOWN(l)[4]:The status of the
peer 10.0.2.2 changed from ESTABLISHED to IDLE. (InstanceName=Public,
StateChangeReason=Hold Timer Expired)
[R3]display gtsm statistics all
GTSM Statistics Table
----------------------------------------------------------------
SlotId  Protocol  Total Counters  Drop Counters  Pass Counters
----------------------------------------------------------------
 0      BGP          83              27             56
 0      OSPF         0               0              0
 0      LDP          0               0              0
----------------------------------------------------------------
```

把R2的配置修改回来，间隔一段时间以后再观察是否有数据包被丢弃掉。

```
[R2-bgp]undo peer 10.0.3.3 ebgp-max-hop
[R2-bgp]peer 10.0.3.3 valid-ttl-hops 1

[R3]display gtsm statistics all
GTSM Statistics Table
----------------------------------------------------------------
SlotId  Protocol  Total Counters  Drop Counters  Pass Counters
----------------------------------------------------------------
 0      BGP          89              27              62
 0      OSPF         0               0               0
 0      LDP          0               0               0
----------------------------------------------------------------
```

这时我们看到已经没有新的数据包被丢弃掉了。

步骤五. BGP 复位

我们在部署路由策略（Route-policy）的时候，经常需要复位BGP对等体关系，而且确保策略的生效。在VRP平台中，有两种重置方式，一种是硬复位(reset)，另一种是软重置（refresh）。

Reset bgp命令会拆除路由器与对等体之间的TCP连接。运行该命令以后，对等体之间的关系会回到IDLE状态，然后经过OPEN、OPENCONFIRM，最后到达ESTABLISHED状态。整个过程会持续30秒以上。

由于路由策略的修改，需要经常复位BGP。但如果采用硬重置将会造成路由器与路由器之间中断达30秒以上。对于核心网络而言，带来的长达30秒的网络

中断是无法容忍的。

Refresh bgp命令不拆除BGP的对等体关系，仅在路由器之间重新同步路由表，这个过程比硬复位短得多。如果路由器中的路由表条目数量不是很庞大，几乎感觉不到网络的中断。

注意上面这两个命令均需要在用户视图下运行。我们首先来观察**Reset bgp**的过程。

在R2上运行**reset bgp all**，并立刻查看BGP对等体关系。

```
<R2>reset bgp all
Dec 21 2011 14:39:26+00:00 R2 %%01BGP/3/STATE_CHG_UPDOWN(l)[3]:The status of the
peer 10.0.1.1 changed from ESTABLISHED to IDLE. (InstanceName=Public,
StateChangeReason=CEASE/Administrative Reset)
Dec 21 2011 14:39:26+00:00 R2 %%01BGP/3/STATE_CHG_UPDOWN(l)[4]:The status of the
peer 10.0.3.3 changed from ESTABLISHED to IDLE. (InstanceName=Public,
StateChangeReason=CEASE/Administrative Reset)
<R2>display bgp peer

 BGP local router ID : 10.0.12.2
 Local AS number : 64512
 Total number of peers : 2              Peers in established state : 0

  Peer            V          AS  MsgRcvd  MsgSent  OutQ  Up/Down       State  PrefRcv

  10.0.1.1        4       64512        0        0     0 00:00:08        Idle       0
  10.0.3.3        4       64513        0        0     0 00:00:08        Idle       0
```

等待约半分钟以后，可以看到路由器之间对等体关系又建立起来了。

```
<R2>
Dec 21 2011 14:39:58+00:00 R2 %%01BGP/3/STATE_CHG_UPDOWN(l)[5]:The status of the
peer 10.0.3.3 changed from OPENCONFIRM to ESTABLISHED. (InstanceName=Public,
StateChangeReason=Up)
<R2>
Dec 21 2011 14:39:58+00:00 R2 %%01BGP/3/STATE_CHG_UPDOWN(l)[6]:The status of the
peer 10.0.1.1 changed from OPENCONFIRM to ESTABLISHED. (InstanceName=Public,
StateChangeReason=Up)
```

上面的命令重置了所有的BGP对等体。有时候我们只希望重置其中一个对等体之间的BGP关系，可以在**Reset bgp**的命令后面加上对等体的IP地址或对等体组的名称。这里，我们仅重置R2和R3之间的对等体关系。

```
<R2>reset bgp 10.0.3.3
Dec 21 2011 14:42:13+00:00 R2 %%01BGP/3/STATE_CHG_UPDOWN(l)[7]:The status of the
```

```
peer 10.0.3.3 changed from ESTABLISHED to IDLE. (InstanceName=Public,
StateChangeReason=CEASE/Administrative Reset)
<R2>display bgp peer

 BGP local router ID : 10.0.12.2
 Local AS number : 64512
 Total number of peers : 2              Peers in established state : 1

  Peer            V          AS  MsgRcvd  MsgSent  OutQ  Up/Down      State PrefRcv

  10.0.1.1        4       64512        5        7     0 00:03:03 Established       0
  10.0.3.3        4       64513        0        0     0 00:00:15        Idle       0
```

可以看到这时候仅有10.0.3.3一个对等体进入了Idle状态，而R2与10.0.1.1仍然保持Established。

为了验证软复位也能更新BGP路由表，我们在R2上增加一个路由策略来观察软重置的效果。

以上步骤中已在R2上发布了它的Loopback 0接口网段到BGP，并且在R3的路由表上已经看到了到达该网段的路由，通过BGP路由表可以看到该路由的Origin属性为IGP。

```
[R3]display bgp routing-table

 BGP Local router ID is 10.0.23.3
 Status codes: * - valid, > - best, d - damped,
               h - history,  i - internal, s - suppressed, S - Stale
               Origin : i - IGP, e - EGP, ? - incomplete


 Total Number of Routes: 1
      Network            NextHop        MED        LocPrf    PrefVal Path/Ogn

      10.0.2.2/32        10.0.2.2        0                     0      100i
```

接下来创建访问控制列表及路由策略将该路由的Origin属性改为EGP。

```
[R2]acl number 2001
[R2-acl-basic-2001]rule permit source 10.0.2.2 0
[R2-acl-basic-2001]quit
[R2]route-policy change_origin permit node 10
Info: New Sequence of this List.
[R2-route-policy]if-match acl 2001
[R2-route-policy]apply origin egp 100
```

将该策略应用在R3上，并软重置R2和R3之间的对等体关系。

```
[R2]bgp 64512
[R2-bgp]peer 10.0.3.3 route-policy change_origin export
[R2-bgp]return
<R2>refresh bgp 10.0.3.3 export
```

立刻查看BGP对等体关系，R2与R3的对等体关系仍然存在。

```
[R2]display bgp peer

 BGP local router ID : 10.0.12.2
 Local AS number : 64512
 Total number of peers : 2              Peers in established state : 2

  Peer         V         AS  MsgRcvd  MsgSent  OutQ  Up/Down       State PrefRcv

  10.0.1.1     4      64512       19       21     0 00:17:28 Established       0
  10.0.3.3     4      64513       16       18     0 00:14:08 Established       0
```

在R3上检查BGP路由表，发现该路由的Origin属性已经变成了EGP。

```
[R3]display bgp routing-table

 BGP Local router ID is 10.0.3.3
 Status codes: * - valid, > - best, d - damped,
               h - history,  i - internal, s - suppressed, S - Stale
               Origin : i - IGP, e - EGP, ? - incomplete


 Total Number of Routes: 1
      Network            NextHop        MED        LocPrf    PrefVal Path/Ogn

      10.0.2.2/32        10.0.2.2       0                     0        100e
```

刚才的**refresh**命令单独使用时仅在出方向（export）上有效，但如果对端设备（R3）想主动更新BGP路由表，还需要在对等体关系上加上参数**keep-all-routes**。

```
[R3-bgp]peer 10.0.2.2 keep-all-routes
```

首先在R2上删除该策略，在R3的入方向上软重置对等体关系，可以看到在R3的BGP路由表上该路由的Origin属性又变回了IGP。

```
[R2]bgp 64512
```

```
[R2-bgp]undo peer 10.0.3.3 route-policy change_origin export

<R3>refresh bgp all import
<R3>display bgp routing-table

 BGP Local router ID is 10.0.3.3
 Status codes: * - valid, > - best, d - damped,
             h - history,  i - internal, s - suppressed, S - Stale
             Origin : i - IGP, e - EGP, ? - incomplete


 Total Number of Routes: 1
     Network          NextHop        MED        LocPrf    PrefVal Path/Ogn

     10.0.2.2/32       10.0.2.2       0                     0        100i
```

附加实验：思考并验证

GTSM主要用于防止哪种类型的攻击？

在联盟环境中是否可以使用fake-as？

最终设备配置

```
<R1>display current-configuration
[V200R001C00SPC200]
#
 sysname R1
#
gtsm default-action drop
#
interface Serial1/0/0
 link-protocol ppp
 ip address 10.0.12.1 255.255.255.0
#
interface LoopBack0
 ip address 10.0.1.1 255.255.255.255
#
bgp 64512
 peer 10.0.2.2 as-number 64512
```

```
 peer 10.0.2.2 connect-interface LoopBack0
 peer 10.0.2.2 password simple huawei
 peer 10.0.2.2 valid-ttl-hops 1
 #
 ipv4-family unicast
  undo synchronization
  peer 10.0.2.2 enable
#
ospf 1 router-id 10.0.1.1
 area 0.0.0.0
  network 10.0.12.1 0.0.0.0
  network 10.0.1.1 0.0.0.0
#
return

<R2>display current-configuration
[V200R001C00SPC200]
#
 sysname R2
#
gtsm default-action drop
#
acl number 2001
 rule 5 permit source 10.0.2.2 0
#
interface Serial1/0/0
 link-protocol ppp
 ip address 10.0.12.2 255.255.255.0
#
interface Serial2/0/0
 link-protocol ppp
 ip address 10.0.23.2 255.255.255.0
#
interface LoopBack0
 ip address 10.0.2.2 255.255.255.255
#
bgp 64512
 peer 10.0.1.1 as-number 64512
 peer 10.0.1.1 connect-interface LoopBack0
 peer 10.0.1.1 password simple huawei
 peer 10.0.1.1 valid-ttl-hops 1
 peer 10.0.3.3 as-number 64513
 peer 10.0.3.3 connect-interface LoopBack0
```

```
 peer 10.0.3.3 fake-as 100
 peer 10.0.3.3 valid-ttl-hops 1
 #
 ipv4-family unicast
  undo synchronization
  network 10.0.2.2 255.255.255.255
  peer 10.0.1.1 enable
  peer 10.0.1.1 next-hop-local
  peer 10.0.3.3 enable
#
ospf 1 router-id 10.0.2.2
 area 0.0.0.0
  network 10.0.12.2 0.0.0.0
  network 10.0.2.2 0.0.0.0
#
route-policy change_origin deny node 10
 if-match acl 2001
 apply origin egp 100
#
 ip route-static 10.0.3.3 255.255.255.255 10.0.23.3
#
return

<R3>display current-configuration
[V200R001C00SPC200]
#
 sysname R3
#
gtsm default-action drop
 gtsm log drop-packet all
#
interface Serial2/0/0
 link-protocol ppp
 ip address 10.0.23.3 255.255.255.0
#
interface LoopBack0
 ip address 10.0.3.3 255.255.255.255
#
bgp 64513
 peer 10.0.2.2 as-number 100
 peer 10.0.2.2 connect-interface LoopBack0
 peer 10.0.2.2 valid-ttl-hops 1
 #
```

```
 ipv4-family unicast
  undo synchronization
  peer 10.0.2.2 enable
#
 ip route-static 10.0.2.2 255.255.255.255 10.0.23.2
#
return
```

实验 3-7 BGP 路由反射器

学习目的

- 掌握BGP反射器的配置方法
- 掌握反射器中三种BGP路由器角色之间信息交流的方法
- 查看反射器发射的路由信息的属性信息
- 掌握Cluster-list的作用及Cluster-id的配置方法

拓扑图

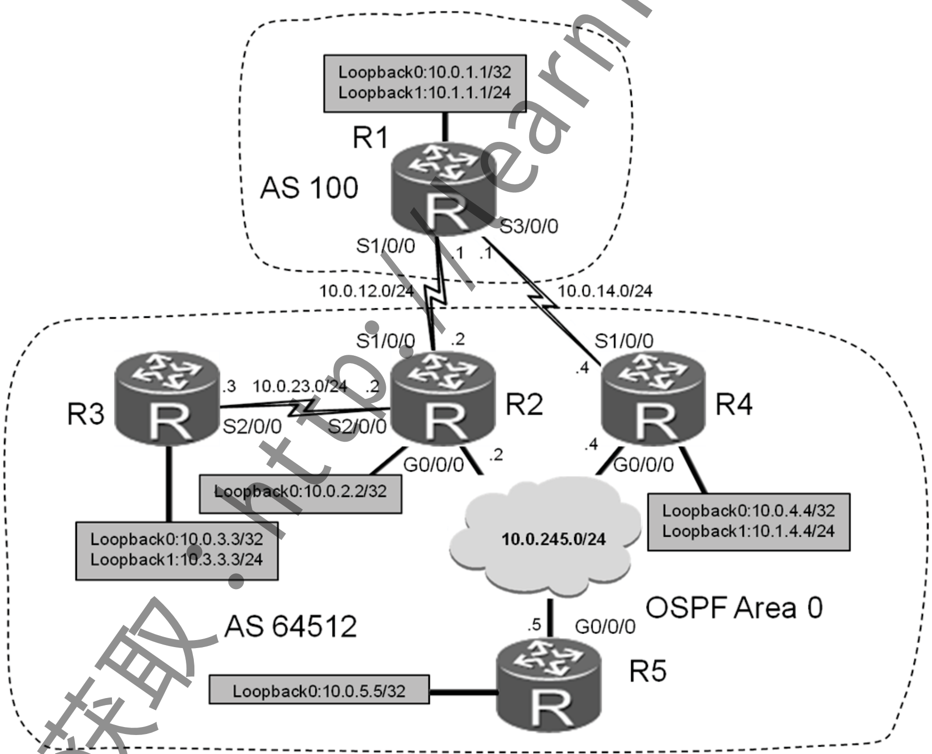

图3-7 BGP路由反射器

场景

你是公司的网络管理员。公司的网络采用了BGP协议与其他AS进行互联，由于BGP本身的特性，要求IBGP对等体之间是全互联的。公司打算在一个AS内部部署路由反射器以减少IBGP对等体关系。在进行详细规划之后，你负责进行了实施工作，最后达到了预期目的。

学习任务

步骤一.　基础配置与 IP 编址

给所有路由器配置物理接口及Loopback接口的IP地址和掩码，各Loopback接口均为32位子网掩码。

```
<R1>system-view
Enter system view, return user view with Ctrl+Z.
[R1]interface Serial 1/0/0
 R1-Serial1/0/0]ip address 10.0.12.1 24
[R1-Serial1/0/0]interface Serial 3/0/0
[R1-Serial3/0/0]ip address 10.0.14.1 24
[R1-Serial3/0/0]interface LoopBack 0
[R1-LoopBack0]ip address 10.0.1.1 32

<R2>system-view
Enter system view, return user view with Ctrl+Z.
[R2]interface Serial 1/0/0
[R2-Serial1/0/0]ip address 10.0.12.2 24
[R2-Serial1/0/0]interface Serial 2/0/0
[R2-Serial2/0/0]ip address 10.0.23.2 24
[R2-Serial2/0/0]interface GigabitEthernet 0/0/0
[R2-GigabitEthernet0/0/0]ip address 10.0.245.2 24
[R2-GigabitEthernet0/0/0]interface LoopBack 0
[R2-LoopBack0]ip address 10.0.2.2 32

<R3>system-view
Enter system view, return user view with Ctrl+Z.
[R3]interface Serial 2/0/0
[R3-Serial2/0/0]ip address 10.0.23.3 24
[R3-Serial2/0/0]interface LoopBack 0
[R3-LoopBack0]ip address 10.0.3.3 32
```

```
<R4>system-view
Enter system view, return user view with Ctrl+Z.
[R4]interface Serial 1/0/0
[R4-Serial1/0/0]ip address 10.0.14.4 24
[R4-Serial1/0/0]interface GigabitEthernet 0/0/0
[R4-GigabitEthernet0/0/0]ip add 10.0.245.4 24
[R4-GigabitEthernet0/0/0]interface LoopBack 0
[R4-LoopBack0]ip address 10.0.4.4 32

<R5>system-view
Enter system view, return user view with Ctrl+Z.
[R5]interface GigabitEthernet 0/0/0
[R5-GigabitEthernet0/0/0]ip address 10.0.245.5 24
[R5-GigabitEthernet0/0/0]interface LoopBack 0
[R5-LoopBack0]ip address 10.0.5.5 32
```

配置完成后，测试直连链路的连通性。

```
[R2]ping -c 1 10.0.12.1
  PING 10.0.12.1: 56  data bytes, press CTRL_C to break
    Reply from 10.0.12.1: bytes=56 Sequence=1 ttl=255 time=41 ms

  --- 10.0.12.1 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 41/41/41 ms

[R2]ping -c 1 10.0.23.3
  PING 10.0.23.3: 56  data bytes, press CTRL_C to break
    Reply from 10.0.23.3: bytes=56 Sequence=1 ttl=255 time=43 ms

  --- 10.0.23.3 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 43/43/43 ms

[R2]ping -c 1 10.0.245.4
  PING 10.0.245.4: 56  data bytes, press CTRL_C to break
    Reply from 10.0.245.4: bytes=56 Sequence=1 ttl=255 time=7 ms

  --- 10.0.245.4 ping statistics ---
    1 packet(s) transmitted
```

```
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 7/7/7 ms

[R2]ping -c 1 10.0.245.5
  PING 10.0.245.5: 56  data bytes, press CTRL_C to break
    Reply from 10.0.245.5: bytes=56 Sequence=1 ttl=255 time=14 ms

  --- 10.0.245.5 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 14/14/14 ms
```

步骤二. 配置 IGP 及 BGP

在AS 64512内部使用OSPF作为IGP，所有设备属于区域0。

R2的G0/0/0、S2/0/0和Loopback 0连接的网段运行OSPF。

```
[R2]ospf 1 router-id 10.0.2.2
[R2-ospf-1]area 0
[R2-ospf-1-area-0.0.0.0]network 10.0.245.2 0.0.0.0
[R2-ospf-1-area-0.0.0.0]network 10.0.23.2 0.0.0.0
[R2-ospf-1-area-0.0.0.0]network 10.0.2.2 0.0.0.0
[R2-ospf-1-area-0.0.0.0]quit
[R2-ospf-1]quit
```

R3的S2/0/0和Loopback 0连接的网段运行OSPF。

```
[R3]ospf 1 router-id 10.0.3.3
[R3-ospf-1]area 0
[R3-ospf-1-area-0.0.0.0]network 10.0.23.3 0.0.0.0
[R3-ospf-1-area-0.0.0.0]network 10.0.3.3 0.0.0.0
[R3-ospf-1-area-0.0.0.0]quit
[R3-ospf-1]quit
```

R4的G0/0/0和Loopback 0连接的网段运行OSPF。

```
[R4]ospf 1 router-id 10.0.4.4
[R4-ospf-1]area 0
[R4-ospf-1-area-0.0.0.0]network 10.0.245.4 0.0.0.0
[R4-ospf-1-area-0.0.0.0]network 10.0.4.4 0.0.0.0
[R4-ospf-1-area-0.0.0.0]quit
```

```
[R4-ospf-1]quit
```

R5的G0/0/0和Loopback 0连接的网段运行OSPF。

```
[R5]ospf 1 router-id 10.0.5.5
[R5-ospf-1]area 0
[R5-ospf-1-area-0.0.0.0]network 10.0.245.5 0.0.0.0
[R5-ospf-1-area-0.0.0.0]network 10.0.5.5 0.0.0.0
[R5-ospf-1-area-0.0.0.0]quit
[R5-ospf-1]quit
```

检查R2已经学习到同区域其他设备的Loopback地址。

```
[R2]display ip routing-table protocol ospf
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Public routing table : OSPF
        Destinations : 3       Routes : 3

OSPF routing table status : <Active>
        Destinations : 3       Routes : 3

Destination/Mask  Proto  Pre  Cost      Flags NextHop       Interface

      10.0.3.3/32  OSPF   10   1562       D   10.0.23.3      Serial2/0/0
      10.0.4.4/32  OSPF   10   1          D   10.0.245.4    GigabitEthernet0/0/0
      10.0.5.5/32  OSPF   10   1          D   10.0.245.5    GigabitEthernet0/0/0

OSPF routing table status : <Inactive>
        Destinations : 0       Routes : 0
```

在R1和R2之间配置EBGP，直接使用物理接口建立对等体连接。

AS的规划如图所示，R1和R4之间的EBGP对等体关系暂不配置。所有的BGP路由器均使用Loopback 0的地址作为Router ID。

```
[R1]bgp 100
[R1-bgp]router-id 10.0.1.1
[R1-bgp]peer 10.0.12.2 as-number 64512

[R2]bgp 64512
[R2-bgp]router-id 10.0.2.2
[R2-bgp]peer 10.0.12.1 as-number 100
```

配置完成以后检查EBGP对等体关系。

```
[R1]display bgp peer

 BGP local router ID : 10.0.12.1
 Local AS number : 100
 Total number of peers : 1              Peers in established state : 1

  Peer          V          AS  MsgRcvd  MsgSent  OutQ  Up/Down       State PrefRcv

  10.0.12.2     4       64512        2        2     0 00:00:24 Established       0
```

步骤三. 配置路由反射器

在这个实验中，我们规划R2作为反射器，R3和R5是R2的客户，R4是非客户。首先在R2上创建对等体组rr_group，用于与R3和R5建立对等体连接。

建立与R4的IBGP对等体关系，为保证AS内路由器能学习到AS 100的路由。

在建立IBGP对等体时加上**next-hop-local**参数。

```
[R2-bgp]group rr_group internal
[R2-bgp]peer 10.0.3.3 group rr_group
[R2-bgp]peer 10.0.5.5 group rr_group
[R2-bgp]peer rr_group connect-interface LoopBack 0
[R2-bgp]peer rr_group reflect-client
[R2-bgp]peer rr_group next-hop-local
[R2-bgp]peer 10.0.4.4 as-number 64512
[R2-bgp]peer 10.0.4.4 connect-interface LoopBack 0
[R2-bgp]peer 10.0.4.4 next-hop-local
```

在R3和R5上分别建立与反射器R2的对等体关系，在反射器的客户上不需要增加额外的配置。

```
[R3]bgp 64512
[R3-bgp]router-id 10.0.3.3
[R3-bgp]peer 10.0.2.2 as-number 64512
[R3-bgp]peer 10.0.2.2 connect-interface LoopBack 0

[R5]bgp 64512
[R5-bgp]router-id 10.0.5.5
[R5-bgp]peer 10.0.2.2 as-number 64512
[R5-bgp]peer 10.0.2.2 connect-interface LoopBack 0
```

在非客户R4上增加到R2的IBGP对等体关系。

```
[R4]bgp 64512
[R4-bgp]router-id 10.0.4.4
[R4-bgp]peer 10.0.2.2 as-number 64512
[R4-bgp]peer 10.0.2.2 connect-interface LoopBack 0
```

在R2上检查BGP对等体均已建立。

```
[R2]display bgp peer

 BGP local router ID : 10.0.12.2
 Local AS number : 64512
 Total number of peers : 4            Peers in established state : 4

  Peer          V          AS  MsgRcvd  MsgSent  OutQ  Up/Down       State PrefRcv

  10.0.3.3      4       64512        4        4     0 00:02:09 Established       0
  10.0.4.4      4       64512        3        4     0 00:01:43 Established       0
  10.0.5.5      4       64512        2        4     0 00:00:09 Established       0
  10.0.12.1     4         100       36       37     0 00:34:31 Established       0
```

在R2上查看BGP对等体的详细信息可以看到R3和R5是R2的客户。

```
[R2]display bgp peer verbose

        BGP Peer is 10.0.3.3,  remote AS 64512
        Type: IBGP link
        BGP version 4, Remote router ID 10.0.23.3
        Update-group ID: 0
        BGP current state: Established, Up for 00h06m02s
        BGP current event: RecvKeepalive
        BGP last state: OpenConfirm
        BGP Peer Up count: 1
        Received total routes: 0
        Received active routes total: 0
        Advertised total routes: 0
        Port:  Local - 179     Remote - 50122
        Configured: Connect-retry Time: 32 sec
        Configured: Active Hold Time: 180 sec   Keepalive Time:60 sec
        Received  : Active Hold Time: 180 sec
        Negotiated: Active Hold Time: 180 sec   Keepalive Time:60 sec
        Peer optional capabilities:
        Peer supports bgp multi-protocol extension
        Peer supports bgp route refresh capability
        Peer supports bgp 4-byte-as capability
```

```
      Address family IPv4 Unicast: advertised and received
 Received: Total 8 messages
              Update messages              0
              Open messages               1
              KeepAlive messages           7
              Notification messages          0
              Refresh messages             0
 Sent: Total 8 messages
              Update messages              0
              Open messages               1
              KeepAlive messages           7
              Notification messages          0
              Refresh messages             0
 Authentication type configured: None
 Last keepalive received: 2011/12/08 15:22:07
 Minimum route advertisement interval is 15 seconds
 Optional capabilities:
 Route refresh capability has been enabled
 4-byte-as capability has been enabled
 It's route-reflector-client
 Connect-interface has been configured
 Peer Preferred Value: 0
 Routing policy configured:
 No routing policy is configured
……output omit……
```

步骤四. 观察反射器中路由信息的传递

为了观察路由信息的传递，我们在路由器上创建额外的Loopback接口，使用**network**命令发布到BGP里。我们首先观察到的情况是IBGP非客户发布路由的场景。

在R4上创建Loopback 1，地址为10.1.4.4/24，并发布到BGP里。

```
[R4]interface LoopBack 1
[R4-LoopBack1]ip address 10.1.4.4 24
[R4-LoopBack1]quit
[R4]bgp 64512
[R4-bgp]network 10.1.4.0 24
```

在R2上检查是否学习到该路由。

```
[R2]display bgp routing-table
```

```
 BGP Local router ID is 10.0.12.2
 Status codes: * - valid, > - best, d - damped,
               h - history,  i - internal, s - suppressed, S - Stale
               Origin : i - IGP, e - EGP, ? - incomplete


 Total Number of Routes: 1
      Network            NextHop        MED        LocPrf    PrefVal Path/Ogn

 *>i  10.1.4.0/24         10.0.4.4       0          100        0      i
```

然后在R3和R5上检查是否存在该路由，R3和R5均是R2的客户，所以均学习到了这条路由。

```
[R3]display bgp routing-table

 BGP Local router ID is 10.0.23.3
 Status codes: * - valid, > - best, d - damped,
               h - history,  i - internal, s - suppressed, S - Stale
               Origin : i - IGP, e - EGP, ? - incomplete


 Total Number of Routes: 1
      Network            NextHop        MED        LocPrf    PrefVal Path/Ogn

 *>i  10.1.4.0/24         10.0.4.4       0          100        0      i

[R5]display bgp routing-table

 BGP Local router ID is 10.0.245.5
 Status codes: * - valid, > - best, d - damped,
               h - history,  i - internal, s - suppressed, S - Stale
               Origin : i - IGP, e - EGP, ? - incomplete


 Total Number of Routes: 1
      Network            NextHop        MED        LocPrf    PrefVal Path/Ogn

 *>i  10.1.4.0/24         10.0.4.4       0          100        0      i
```

这里我们看出了反射器通告路由的第一条原则：**从非客户机IBGP对等体学到的路由，发布给此反射器的所有客户机。**

其实这时候我们可以注意到R1也学习到了这条路由。这并不是因为R2是反射器，才把10.1.4.0/24通告给R1。而是因为BGP路由器从IBGP对等体学习到的路由仅不通告给其他IBGP对等体，但会通告给EBGP对等体。

考虑路由信息由客户发起的情况。

在R3上创建Loopback 1，地址为10.1.3.3/24，将该网段发布到BGP里。

```
[R3]interface LoopBack 1
[R3-LoopBack1]ip address 10.1.3.3 24
[R3-LoopBack1]bgp 64512
[R3-bgp]network 10.1.3.0 24
[R3-bgp]quit
```

在R4和R5上分别查看路由表。

```
[R4]display bgp routing-table

 BGP Local router ID is 10.0.14.4
 Status codes: * - valid, > - best, d - damped,
               h - history,  i - internal, s - suppressed, S - Stale
               Origin : i - IGP, e - EGP, ? - incomplete


 Total Number of Routes: 2
      Network            NextHop        MED        LocPrf    PrefVal Path/Ogn

 *>i  10.1.3.0/24         10.0.3.3       0          100        0      i
 *>   10.1.4.0/24        0.0.0.0         0                     0      i

[R5]display bgp routing-table

 BGP Local router ID is 10.0.245.5
 Status codes: * - valid, > - best, d - damped,
               h - history,  i - internal, s - suppressed, S - Stale
               Origin : i - IGP, e - EGP, ? - incomplete


 Total Number of Routes: 2
      Network            NextHop        MED        LocPrf    PrefVal Path/Ogn

 *>i  10.1.3.0/24         10.0.3.3       0          100        0      i
 *>i  10.1.4.0/24         10.0.4.4       0          100        0      i
```

我们可以看到在R4和R5上均学习到了这条路由，这是反射器通告路由的第

二条原则：**从客户机学到的路由，发布给此反射器的所有非客户机和客户机（发起此路由的客户机除外）。**

下面考虑路由更新由EBGP对等体发起的情况。

我们在R1上创建Loopback1，地址为10.1.1.1/24，并发布到BGP里。

```
[R1]interface LoopBack 1
[R1-LoopBack1]ip address 10.1.1.1 24
[R1-LoopBack1]bgp 100
[R1-bgp]network 10.1.1.0 24
```

在R4和R5上分别查看路由表。

```
[R4]display bgp routing-table

 BGP Local router ID is 10.0.14.4
 Status codes: * - valid, > - best, d - damped,
               h - history,  i - internal, s - suppressed, S - Stale
               Origin : i - IGP, e - EGP, ? - incomplete


 Total Number of Routes: 3
      Network            NextHop        MED        LocPrf    PrefVal Path/Ogn

 *>i  10.1.1.0/24         10.0.2.2       0          100        0      100i
 *>i  10.1.3.0/24         10.0.3.3       0          100        0      i
 *>   10.1.4.0/24         0.0.0.0         0                     0      i

[R5]display bgp routing-table

 BGP Local router ID is 10.0.245.5
 Status codes: * - valid, > - best, d - damped,
               h - history,  i - internal, s - suppressed, S - Stale
               Origin : i - IGP, e - EGP, ? - incomplete


 Total Number of Routes: 3
      Network            NextHop        MED        LocPrf    PrefVal Path/Ogn

 *>i  10.1.1.0/24         10.0.2.2       0          100        0      100i
 *>i  10.1.3.0/24         10.0.3.3       0          100        0      i
 *>i  10.1.4.0/24         10.0.4.4       0          100        0      i
```

这里可以看到在R4和R5上均学习到了这条路由，这是反射器通告路由的第

三条原则：**从EBGP对等体学到的路由，发布给所有的非客户机和客户机。**

步骤五. 查看反射器发射的路由信息的属性

我们首先来研究在有反射器的场景中的Next-hop属性。

查看R3的BGP路由表。

```
[R3]display bgp routing-table

 BGP Local router ID is 10.0.3.3
 Status codes: * - valid, > - best, d - damped,
               h - history,  i - internal, s - suppressed, S - Stale
               Origin : i - IGP, e - EGP, ? - incomplete


 Total Number of Routes: 3
      Network            NextHop        MED      LocPrf    PrefVal  Path/Ogn

 *>i  10.1.1.0/24        10.0.2.2        0         100        0      100i
 *>   10.1.3.0/24        0.0.0.0         0                     0      i
 *>i  10.1.4.0/24        10.0.4.4        0         100        0      i
```

注意到在R3上共有2条从其他路由器上学习来的路由，其中10.1.1.0/24为R1发布的路由，是R2从EBGP对等体这里学习到的。R2上配置了**Next-hop-local**，所以R3学习到这条路由的时候下一跳是R2的Loopback地址。

10.1.4.0/24是R3从R4这里学习来的。R2反射了这条路由，我们可以看到这条路由的下一跳还是10.0.4.4，也就是这条路由的始发者。

从这里可以看出，虽然在R2上配置了**Next-hop-local**，但反射器没有改变非客户通告的路由的下一跳直接通告给客户。

下面观察反射器将客户的路由通告给非客户的场景。

首先在R5上配置Loopback1，并将该网段地址通过BGP通告出去。

```
[R5]interface LoopBack 1
[R5-LoopBack1]ip address 10.1.5.5 24
[R5]bgp 64512
[R5-bgp]network 10.1.5.0 24
```

在R4上观察R5通告的路由10.1.5.0/24。

```
[R4]display bgp routing-table
```

```
 BGP Local router ID is 10.0.4.4
 Status codes: * - valid, > - best, d - damped,
               h - history,  i - internal, s - suppressed, S - Stale
               Origin : i - IGP, e - EGP, ? - incomplete


 Total Number of Routes: 4
     Network           NextHop        MED        LocPrf    PrefVal Path/Ogn

 *>i  10.1.1.0/24       10.0.2.2       0          100        0      100i
 *>i  10.1.3.0/24       10.0.3.3       0          100        0      i
 *>   10.1.4.0/24       0.0.0.0        0                     0      i
 *>i  10.1.5.0/24       10.0.5.5       0          100        0      i
```

可以看到该路由的下一跳是R5的环回接口地址。

得出一个结论：反射器不会改变从IBGP对等体那里学习到的路由信息的下一跳。

接下来讨论Originator_ID属性和Cluster_List属性。

由于路由反射器打破了BGP路由器从IBGP对等体学习到的路由不通告给其他对等体的规定，所以就有可能在AS内产生环路。Originator_ID和Cluster_List就是用于检测和防止路由环路的。在R4上查看BGP路由10.1.4.0的详细信息，可以看到这两个属性的值。

```
[R3]display bgp routing-table 10.1.4.0

 BGP local router ID : 10.0.3.3
 Local AS number : 64512
 Paths:   1 available, 1 best, 1 select
 BGP routing table entry information of 10.1.4.0/24:
 From: 10.0.2.2 (10.0.2.2)
 Route Duration: 00h02m07s
 Relay IP Nexthop: 10.0.23.2
 Relay IP Out-Interface: Serial2/0/0
 Original nexthop: 10.0.4.4
 Qos information : 0x0
 AS-path Nil, origin igp, MED 0, localpref 100, pref-val 0, valid, internal, best,
select, active, pre 255
 Originator:  10.0.4.4
 Cluster list: 10.0.2.2
 Not advertised to any peer yet
```

当一条路由第一次被反射器反射的时候，反射器将Originator属性加入这条

路由，标识这条路由的发起路由器。

如果一条路由中已经存在了Originator属性，则反射器将不会创建新的Originator。

10.1.4.0这条路由的始发者是10.0.4.4，所以反射器R2将该值填入路由的属性中。当BGP路由器接收到路由的时候，将比较收到的Originator和本地的Router ID。如果两个ID相同，BGP 路由器会忽略掉这条路由，不做处理，以避免出现环路。

Cluster list描述了一条路由所经过的反射器的列表，这里10.1.4.0这条路由经过了R2这一台反射器，所以在Cluster list看到了10.0.2.2。

步骤六. 配置 Cluster ID

为增加网络的可靠性，防止单点故障，有时需要在一个集群中配置一个以上的路由反射器。这时，相同集群中的路由反射器应该配置相同的Cluster ID，以避免路由环路。

在这个实验中，将R4配置成一台路由反射器，R5作为R4的客户。同时配置R1和R4之间的EBGP对等体关系。

```
[R1]bgp 100
[R1-bgp]peer 10.0.14.4 as-number 64512

[R4]bgp 64512
[R4-bgp]peer 10.0.14.1 as-number 100
[R4-bgp]group rr_group internal
[R4-bgp]peer 10.0.5.5 group rr_group
[R4-bgp]peer rr_group connect-interface LoopBack 0
[R4-bgp]peer rr_group next-hop-local
[R4-bgp]peer rr_group reflect-client

[R5]bgp 64512
[R5-bgp]peer 10.0.4.4 as-number 64512
[R5-bgp]peer 10.0.4.4 connect-interface LoopBack 0
```

这样，一共有2台反射器对R5提供路由的反射。当反射器接收到一条更新路由时，它会检查Cluster_List。如果Cluster_List中已经有本地Cluster_ID，就丢弃该路由。我们在这里必须把R2和R4的Cluster ID改成一样，这样就防止了路由环路。

```
[R2-bgp]reflector cluster-id 1

[R4-bgp]reflector cluster-id 1
```

为了让R5能同时选中到达10.1.1.0/24的两条路径，在R5上将BGP负载分担的条数改成2。

```
[R5-bgp]maximum load-balancing 2
```

这时在R5上查看BGP路由的详细信息。首先查看路由10.1.4.0。

```
[R5]display bgp routing-table 10.1.4.0

 BGP local router ID : 10.0.5.5
 Local AS number : 64512
 Paths:   2 available, 1 best, 1 select
 BGP routing table entry information of 10.1.4.0/24:
 From: 10.0.4.4 (10.0.4.4)
 Route Duration: 00h34m36s
 Relay IP Nexthop: 10.0.245.4
 Relay IP Out-Interface: GigabitEthernet0/0/0
 Original nexthop: 10.0.4.4
 Qos information : 0x0
 AS-path Nil, origin igp, MED 0, localpref 100, pref-val 0, valid, internal, best,
select, active, pre 255
 Not advertised to any peer yet

 BGP routing table entry information of 10.1.4.0/24:
 From: 10.0.2.2 (10.0.2.2)
 Route Duration: 00h00m51s
 Relay IP Nexthop: 10.0.245.4
 Relay IP Out-Interface: GigabitEthernet0/0/0
 Original nexthop: 10.0.4.4
 Qos information : 0x0
 AS-path Nil, origin igp, MED 0, localpref 100, pref-val 0, valid, internal, pre
255, not preferred for Cluster List
 Originator:  10.0.4.4
 Cluster list: 0.0.0.1
 Not advertised to any peer yet
```

可以看到R5从两台路由器上学习到了10.1.4.0/24这条路由，其中Router-ID为10.0.4.4的路由器是这条路由的始发者，它发布的路由信息中是没有Originator和Cluster list属性。

R2（Router-ID为10.0.2.2）是反射器。R2将该路由再反射给R5，并且将Originator设为路由的始发者10.0.4.4，在Cluster list中添加上自己的Cluster ID。

由于R4是直接向R5传递路由信息的，而R2则经过一次反射以后把路由通告给R5，所以R5优选R4通告的BGP路由。

```
[R5]display bgp routing-table 10.1.3.0

 BGP local router ID : 10.0.5.5
 Local AS number : 64512
 Paths:   1 available, 1 best, 1 select
 BGP routing table entry information of 10.1.3.0/24:
 From: 10.0.2.2 (10.0.2.2)
 Route Duration: 00h01m17s
 Relay IP Nexthop: 10.0.245.2
 Relay IP Out-Interface: GigabitEthernet0/0/0
 Original nexthop: 10.0.3.3
 Qos information : 0x0
 AS-path Nil, origin igp, MED 0, localpref 100, pref-val 0, valid, internal, best,
select, active, pre 255
 Originator:  10.0.3.3
 Cluster list: 0.0.0.1
 Not advertised to any peer yet
```

注意到10.1.3.0/24这条路由仅从R2这里学习到，反射器R4并没有将这条路由再次发送给R5，接下来在R4上查看10.1.3.0/24这条路由。

```
[R4]display bgp routing-table

 BGP Local router ID is 10.0.4.4
 Status codes: * - valid, > - best, d - damped,
               h - history,  i - internal, s - suppressed, S - Stale
               Origin : i - IGP, e - EGP, ? - incomplete


 Total Number of Routes: 3
     Network          NextHop        MED        LocPrf    PrefVal Path/Ogn

 *>   10.1.1.0/24    10.0.14.1       0                     0      100i
 * I                 10.0.2.2        0          100        0      100i
 *>   10.1.4.0/24    0.0.0.0         0                     0      i
 *>i  10.1.5.0/24    10.0.245.5      0          100        0      i
```

可以发现R4并没有学习到10.1.3.0/24这条路由。分析原因可知，R3是R2的客户，因此R3发布的路由会被R2在Cluster list里加上自己的Cluster ID：0.0.0.1。由于R2和R4之间使用了相同的Cluster ID，R4会丢弃R2发来的10.1.3.0/24这条路由。

由此可知在构建一个反射器的集群以后，集群内的客户机必须和所有的反射器有IBGP对等体关系。

在R3和R4之间创建对等体关系。

```
[R4-bgp]peer 10.0.3.3 group rr_group

[R3-bgp]peer 10.0.4.4 as-number 64512
[R3-bgp]peer 10.0.4.4 connect-interface LoopBack 0
```

检查R4的路由表。

```
[R4]display bgp routing-table

 BGP Local router ID is 10.0.4.4
 Status codes: * - valid, > - best, d - damped,
               h - history,  i - internal, s - suppressed, S - Stale
               Origin : i - IGP, e - EGP, ? - incomplete


 Total Number of Routes: 4
     Network          NextHop       MED      LocPrf    PrefVal Path/Ogn

 *>   10.1.1.0/24   10.0.14.1       0                    0      100i
 * I                10.0.2.2        0         100        0      100i
 *>i  10.1.3.0/24   10.0.3.3        0         100        0      i
 *>   10.1.4.0/24   0.0.0.0         0                    0      i
 *>i  10.1.5.0/24   10.0.245.5      0         100        0      i
```

可以看到10.1.3.0已经出现了。

在R5上检查10.1.3.0这条路由的详细信息，可以看到R5同时从R2和R4这里收到了这条路由，经过反射以后，Cluster list里加上了值为0.0.0.1的Cluster ID。

```
[R5]display bgp routing-table 10.1.3.0

 BGP local router ID : 10.0.5.5
 Local AS number : 64512
 Paths:   2 available, 1 best, 1 select
 BGP routing table entry information of 10.1.3.0/24:
 From: 10.0.2.2 (10.0.2.2)
 Route Duration: 16h16m56s
 Relay IP Nexthop: 10.0.245.2
 Relay IP Out-Interface: GigabitEthernet0/0/0
 Original nexthop: 10.0.3.3
 Qos information : 0x0
```

```
 AS-path Nil, origin igp, MED 0, localpref 100, pref-val 0, valid, internal, best,
select, active, pre 255
 Originator:  10.0.3.3
 Cluster list: 0.0.0.1
 Not advertised to any peer yet

 BGP routing table entry information of 10.1.3.0/24:
 From: 10.0.4.4 (10.0.4.4)
 Route Duration: 00h07m25s
 Relay IP Nexthop: 10.0.245.2
 Relay IP Out-Interface: GigabitEthernet0/0/0
 Original nexthop: 10.0.3.3
 Qos information : 0x0
 AS-path Nil, origin igp, MED 0, localpref 100, pref-val 0, valid, internal, pre
255, not preferred for peer address
 Originator:  10.0.3.3
 Cluster list: 0.0.0.1
 Not advertised to any peer yet
```

附加实验：思考并验证

思考一下，在同一个集群里，反射器是否可以互为客户？

最终设备配置

```
<R1>display current-configuration
[V200R001C00SPC200]
#
 sysname R1
#
interface Serial1/0/0
 link-protocol ppp
 ip address 10.0.12.1 255.255.255.0
#
interface Serial3/0/0
 link-protocol ppp
 ip address 10.0.14.1 255.255.255.0
#
interface LoopBack0
 ip address 10.0.1.1 255.255.255.255
```

```
#
interface LoopBack1
 ip address 10.1.1.1 255.255.255.0
#
bgp 100
 router-id 10.0.1.1
 peer 10.0.12.2 as-number 64512
 peer 10.0.14.4 as-number 64512
 #
 ipv4-family unicast
  undo synchronization
  network 10.1.1.0 255.255.255.0
  peer 10.0.12.2 enable
  peer 10.0.14.4 enable
#
return
```

```
<R2>display current-configuration
[V200R001C00SPC200]
#
 sysname R1
#
interface Serial1/0/0
 link-protocol ppp
 ip address 10.0.12.2 255.255.255.0
#
interface Serial2/0/0
 link-protocol ppp
 ip address 10.0.23.2 255.255.255.0
#
interface GigabitEthernet0/0/0
 ip address 10.0.245.2 255.255.255.0
#
interface LoopBack0
 ip address 10.0.2.2 255.255.255.255
#
bgp 64512
 router-id 10.0.2.2
 peer 10.0.4.4 as-number 64512
 peer 10.0.4.4 connect-interface LoopBack0
 peer 10.0.12.1 as-number 100
 group rr_group internal
 peer rr_group connect-interface LoopBack0
```

```
 peer 10.0.3.3 as-number 64512
 peer 10.0.3.3 group rr_group
 peer 10.0.5.5 as-number 64512
 peer 10.0.5.5 group rr_group
 #
 ipv4-family unicast
  undo synchronization
  reflector cluster-id 1
  peer 10.0.4.4 enable
  peer 10.0.4.4 next-hop-local
  peer 10.0.12.1 enable
  peer rr_group enable
  peer rr_group reflect-client
  peer rr_group next-hop-local
  peer 10.0.3.3 enable
  peer 10.0.3.3 group rr_group
  peer 10.0.5.5 enable
  peer 10.0.5.5 group rr_group
#
ospf 1 router-id 10.0.2.2
 area 0.0.0.0
  network 10.0.12.2 0.0.0.0
  network 10.0.2.2 0.0.0.0
  network 10.0.245.2 0.0.0.0
  network 10.0.23.2 0.0.0.0
#
return

<R3>display current-configuration
[V200R001C00SPC200]
#
 sysname R1
#
interface Serial2/0/0
 link-protocol ppp
 ip address 10.0.23.3 255.255.255.0
#
interface LoopBack0
 ip address 10.0.3.3 255.255.255.255
#
interface LoopBack1
 ip address 10.1.3.3 255.255.255.0
#
```

```
bgp 64512
 router-id 10.0.3.3
 peer 10.0.2.2 as-number 64512
 peer 10.0.2.2 connect-interface LoopBack0
 peer 10.0.4.4 as-number 64512
 peer 10.0.4.4 connect-interface LoopBack0
 #
 ipv4-family unicast
  undo synchronization
  network 10.1.3.0 255.255.255.0
  peer 10.0.2.2 enable
  peer 10.0.4.4 enable
#
ospf 1 router-id 10.0.3.3
 area 0.0.0.0
  network 10.0.23.3 0.0.0.0
  network 10.0.3.3 0.0.0.0
#
return

<R4>display current-configuration
[V200R001C00SPC200]
#
 sysname R1
#
interface Serial1/0/0
 link-protocol ppp
 ip address 10.0.14.4 255.255.255.0
#
interface GigabitEthernet0/0/0
 ip address 10.0.245.4 255.255.255.0
#
interface LoopBack0
 ip address 10.0.4.4 255.255.255.255
#
interface LoopBack1
 ip address 10.1.4.4 255.255.255.0
#
bgp 64512
 router-id 10.0.4.4
 peer 10.0.2.2 as-number 64512
 peer 10.0.2.2 connect-interface LoopBack0
 peer 10.0.14.1 as-number 100
```

```
 group rr_group internal
 peer rr_group connect-interface LoopBack0
 peer 10.0.3.3 as-number 64512
 peer 10.0.3.3 group rr_group
 peer 10.0.5.5 as-number 64512
 peer 10.0.5.5 group rr_group
 #
 ipv4-family unicast
  undo synchronization
  reflector cluster-id 1
  network 10.1.4.0 255.255.255.0
  peer 10.0.2.2 enable
  peer 10.0.14.1 enable
  peer rr_group enable
  peer rr_group reflect-client
  peer rr_group next-hop-local
  peer 10.0.3.3 enable
  peer 10.0.3.3 group rr_group
  peer 10.0.5.5 enable
  peer 10.0.5.5 group rr_group
#
ospf 1 router-id 10.0.4.4
 area 0.0.0.0
  network 10.0.245.4 0.0.0.0
  network 10.0.4.4 0.0.0.0
#
return


<R5>display current-configuration
[V200R001C00SPC200]
#
 sysname R1
#
interface GigabitEthernet0/0/0
 ip address 10.0.245.5 255.255.255.0
#
interface LoopBack0
 ip address 10.0.5.5 255.255.255.255
#
interface LoopBack1
 ip address 10.1.5.5 255.255.255.0
#
```

```
bgp 64512
 router-id 10.0.5.5
 peer 10.0.2.2 as-number 64512
 peer 10.0.2.2 connect-interface LoopBack0
 peer 10.0.4.4 as-number 64512
 peer 10.0.4.4 connect-interface LoopBack0
 #
 ipv4-family unicast
  undo synchronization
  network 10.1.4.0 255.255.255.0
  maximum load-balancing 2
  peer 10.0.2.2 enable
  peer 10.0.4.4 enable
#
ospf 1 router-id 10.0.5.5
 area 0.0.0.0
  network 10.0.245.5 0.0.0.0
  network 10.0.5.5 0.0.0.0
#
return
```

实验 3-8 BGP 联盟（选做）

学习目的

- 掌握隐藏私有BGP AS号的方法
- 掌握联盟的配置方法
- 理解BGP联盟环境中各种BGP属性的变化情况

拓扑图

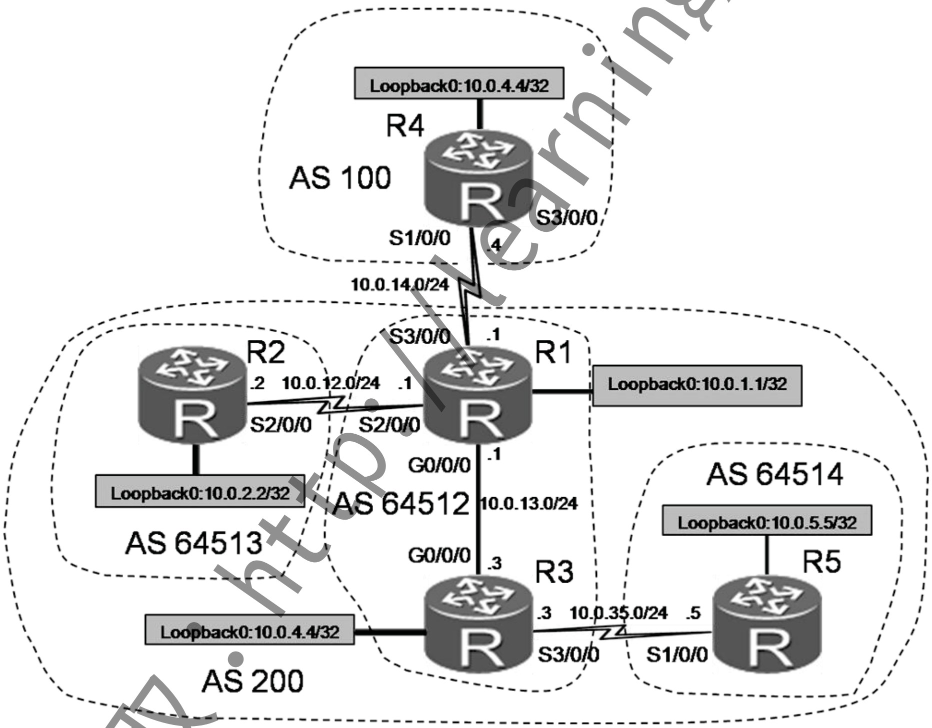

图3-8 BGP联盟

场景

你是公司的网络管理员。公司向运营商申请了一个公网的AS号200，并接入了Internet。随着公司的发展，一个AS已不能满足网络发展的需要，但由于公网

AS号比较紧张，公司只能使用私有的AS号。为了保证私有AS之内的发布的路由也能被公网上其他路由器学习到，你首先采用了隐藏私有AS号的方法接入Internet。随后公司又部署了联盟，这样，公司所有的路由器均部署在了AS200中，每个分支机构使用了一个私有的AS号。

学习任务

步骤一. 基础配置与 IP 编址

给所有路由器配置物理接口及Loopback接口的IP地址和掩码，注意各Loopback0接口地址均为32位掩码。

```
<R1>system-view
Enter system view, return user view with Ctrl+Z.
[R1]interface Serial 1/0/0
[R1-Serial1/0/0]ip address 10.0.12.1 24
[R1-Serial1/0/0]interface Serial 3/0/0
[R1-Serial3/0/0]ip address 10.0.14.1 24
[R1-Serial3/0/0]interface GigabitEthernet 0/0/0
[R1-GigabitEthernet0/0/0]ip address 10.0.13.1 24
[R1-GigabitEthernet0/0/0]interface LoopBack 0
[R1-LoopBack0]ip add 10.0.1.1 32

<R2>system-view
Enter system view, return user view with Ctrl+Z.
[R2]interface Serial 1/0/0
[R2-Serial1/0/0]ip address 10.0.12.2 24
[R2-Serial1/0/0]interface LoopBack 0
[R2-LoopBack0]ip address 10.0.2.2 32

<R3>system-view
Enter system view, return user view with Ctrl+Z.
[R3]interface GigabitEthernet 0/0/0
[R3-GigabitEthernet0/0/0]ip address 10.0.13.3 24
[R3-GigabitEthernet0/0/0]interface Serial 3/0/0
[R3-Serial3/0/0]ip address 10.0.35.3 24
[R3-Serial3/0/0]interface LoopBack 0
[R3-LoopBack0]ip address 10.0.3.3 32

<R4>system-view
Enter system view, return user view with Ctrl+Z.
[R4]interface Serial 1/0/0
```

```
[R4-Serial1/0/0]ip address 10.0.14.4 24
[R4-Serial1/0/0]interface GigabitEthernet 0/0/0
[R4-GigabitEthernet0/0/0]ip add 10.0.245.4 24
[R4-Serial1/0/0]interface LoopBack 0
[R4-LoopBack0]ip address 10.0.4.4 32

<R5>system-view
Enter system view, return user view with Ctrl+Z.
[R5]interface Serial 1/0/0
[R5-Serial1/0/0]ip address 10.0.35.5 24
[R5-Serial1/0/0]interface LoopBack 0
[R5-LoopBack0]ip address 10.0.5.5 32
```

配置完成后，测试直连链路的连通性。

```
[R1]ping -c 1 10.0.14.4
  PING 10.0.14.4: 56  data bytes, press CTRL_C to break
    Reply from 10.0.14.4: bytes=56 Sequence=1 ttl=255 time=41 ms

  --- 10.0.14.4 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 41/41/41 ms

[R1]ping -c 1 10.0.12.2
  PING 10.0.12.2: 56  data bytes, press CTRL_C to break
    Reply from 10.0.12.2: bytes=56 Sequence=1 ttl=255 time=34 ms

  --- 10.0.12.2 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 34/34/34 ms

[R1]ping -c 1 10.0.13.3
  PING 10.0.13.3: 56  data bytes, press CTRL_C to break
    Reply from 10.0.13.3: bytes=56 Sequence=1 ttl=255 time=13 ms

  --- 10.0.13.3 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 13/13/13 ms
```

```
[R5]ping -c 1 10.0.35.3
  PING 10.0.35.3: 56 data bytes, press CTRL_C to break
    Reply from 10.0.35.3: bytes=56 Sequence=1 ttl=255 time=36 ms

  --- 10.0.35.3 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 36/36/36 ms
```

步骤二. 配置 BGP 连接

公司最早没有部署联盟，采用BGP直接连接到运营商且只有AS 200一个自治系统。随着公司的发展，一个自治系统已不能满足公司网络发展的需要，因此公司有增加了其他自治系统，但使用的是私有的自治系统号。

所有的BGP对等体关系均采用物理接口直接建立对等体关系。

AS的规划如拓扑图所示。

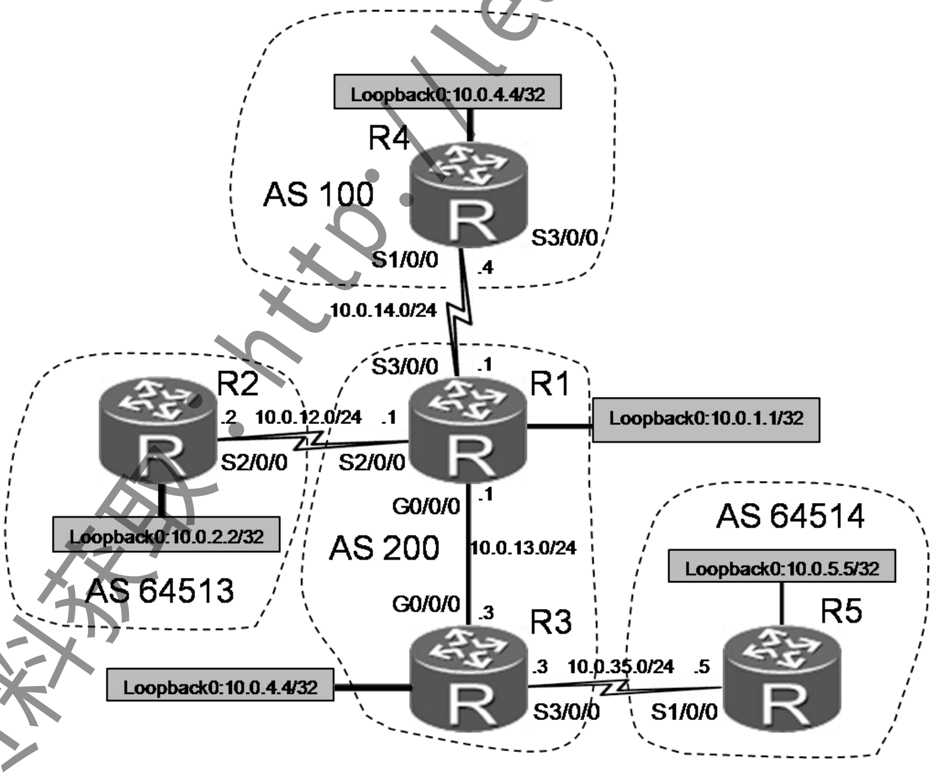

```
[R1]bgp 200
[R1-bgp]router-id 10.0.1.1
[R1-bgp]peer 10.0.14.4 as-number 100
[R1-bgp]peer 10.0.13.3 as-number 200
[R1-bgp]peer 10.0.13.3 next-hop-local
[R1-bgp]peer 10.0.12.2 as-number 64513

[R2]bgp 64513
[R2-bgp]router-id 10.0.2.2
[R2-bgp]peer 10.0.12.1 as-number 200

[R3]bgp 200
[R3-bgp]router-id 10.0.3.3
[R3-bgp]peer 10.0.13.1 as-number 200
[R3-bgp]peer 10.0.13.1 next-hop-local
[R3-bgp]peer 10.0.35.5 as-number 64514

[R4]bgp 100
[R4-bgp]router-id 10.0.4.4
[R4-bgp]peer 10.0.14.1 as-number 200

[R5]bgp 64514
[R5-bgp]router-id 10.0.5.5
[R5-bgp]peer 10.0.35.3 as-number 200
```

配置完成以后在R1上检查对等体关系。

```
[R1]display bgp peer

 BGP local router ID : 10.0.1.1
 Local AS number : 64512
 Total number of peers : 3           Peers in established state : 2

  Peer            V          AS  MsgRcvd  MsgSent  OutQ  Up/Down       State PrefRcv

  10.0.12.2       4       64513       17       20     0 00:15:24 Established       1
  10.0.13.3       4       64512        6        7     0 00:06:02 Established       1
  10.0.14.4       4         100        6        5     0 00:04:44 Established       1
```

步骤三. 配置路由发布及 AS 隐藏

为了方便观察路由信息在AS之间的传递，我们将路由器各自的Loopback接口连接的网络通过**network**命令发布到BGP中。

```
[R1]bgp 200
[R1-bgp]network 10.0.1.1 32

[R2]bgp 64513
[R2-bgp]network 10.0.2.2 32

[R3]bgp 200
[R3-bgp]network 10.0.3.3 32

[R4]bgp 100
[R4-bgp]network 10.0.4.4 32

[R5]bgp 64514
[R5-bgp]network 10.0.5.5 32
```

当完成路由发布以后，在R4上检查BGP路由表。

```
[R4]display bgp routing-table

 BGP Local router ID is 10.0.4.4
 Status codes: * - valid, > - best, d - damped,
               h - history,  i - internal, s - suppressed, S - Stale
               Origin : i - IGP, e - EGP, ? - incomplete


 Total Number of Routes: 5
     Network            NextHop       MED       LocPrf    PrefVal Path/Ogn

 *>  10.0.1.1/32        10.0.14.1      0                    0      200i
 *>  10.0.2.2/32        10.0.14.1                           0      200 64513i
 *>  10.0.3.3/32        10.0.14.1                           0      200i
 *>  10.0.4.4/32        0.0.0.0       0                     0      i
 *>  10.0.5.5/32        10.0.14.1                           0      200 64514i
```

在这里我们可以观察到对于10.0.2.2/32和10.0.5.5/32。

由于路由是在AS 64513和AS 64514始发的，所以在路由的AS-Path属性里含有这2个AS号。

公司只申请了一个公网的AS号，含有私有AS号的路由信息不可以在公网上传递，所以我们必须在R1上将私有AS号去掉。

使用**public-as-only**这条命令可实现该功能。

在R1上对R4配置仅发送公有AS信息。

```
[R1-bgp]peer 10.0.14.4 public-as-only
```

配置该命令以后，我们在R4上检查BGP路由表是否发生了改变。

```
[R4]display bgp routing-table

 BGP Local router ID is 10.0.14.4
 Status codes: * - valid, > - best, d - damped,
               h - history,  i - internal, s - suppressed, S - Stale
               Origin : i - IGP, e - EGP, ? - incomplete


 Total Number of Routes: 5
     Network            NextHop        MED        LocPrf    PrefVal Path/Ogn

 *>  10.0.1.1/32         10.0.14.1      0                     0      200i
 *>  10.0.2.2/32         10.0.14.1                            0      200i
 *>  10.0.3.3/32         10.0.14.1                            0      200i
 *>  10.0.4.4/32         0.0.0.0       0                     0      i
 *>  10.0.5.5/32         10.0.14.1                            0      200i
```

这时我们可以看到，10.0.2.2/32和10.0.5.5/32这两条路由的AS-Path属性中已经没有了私有AS号。

下面测试从R4到R2、R5的Loopback0接口地址的连通性。

由于路由器之间的互联地址没有被发布进BGP，所以我们采用带源地址的扩展**ping**，测试网络连通性。

```
[R4]ping -c 1 -a 10.0.4.4 10.0.2.2
  PING 10.0.2.2: 56  data bytes, press CTRL_C to break
    Reply from 10.0.2.2: bytes=56 Sequence=1 ttl=254 time=70 ms

  --- 10.0.2.2 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 70/70/70 ms

[R4]ping -c 1 -a 10.0.4.4 10.0.5.5
```

```
PING 10.0.5.5: 56  data bytes, press CTRL_C to break
  Reply from 10.0.5.5: bytes=56 Sequence=1 ttl=253 time=71 ms

--- 10.0.5.5 ping statistics ---
  1 packet(s) transmitted
  1 packet(s) received
  0.00% packet loss
  round-trip min/avg/max = 71/71/71 ms
```

这样，我们就完成了公司网络到外部网络的连通性。

步骤四.　配置联盟

公司近期计划部署BGP联盟来代替原来在R1上隐藏私有AS号的方式接入Internet。初步的规划是将整个公司的网络配置到AS 200中，原来AS200的路由器使用私有AS号64512。

配置联盟将需要重新配置BGP。

在R1和R3上删除原有的BGP进程，使用新的AS号启动BGP进程，然后配置对等体关系。

```
[R1]undo bgp
Warning: All BGP configurations will be deleted. Continue? [Y/N]: y
[R1]bgp 64512
[R1-bgp]router-id 10.0.1.1
[R1-bgp]confederation id 200
[R1-bgp]confederation peer-as 64513
[R1-bgp]peer 10.0.14.4 as-number 100
[R1-bgp]peer 10.0.13.3 as-number 64512
[R1-bgp]peer 10.0.13.3 next-hop-local
[R1-bgp]peer 10.0.12.2 as-number 64513
[R1-bgp]network 10.0.1.1 32

[R3]undo bgp
Warning: All BGP configurations will be deleted. Continue? [Y/N]: y
[R3]bgp 64512
[R3-bgp]confederation id 200
[R3-bgp]confederation peer-as 64514
[R3-bgp]peer 10.0.13.1 as-number 64512
[R3-bgp]peer 10.0.13.1 next-hop-local
[R3-bgp]peer 10.0.35.5 as-number 64514
[R3-bgp]network 10.0.3.3 32
```

接下来在R2和R5上修改对等体关系，并加上联盟的配置。

```
[R2]bgp 64513
[R2-bgp]undo peer 10.0.12.1
[R2-bgp]confederation id 200
[R2-bgp]confederation peer-as 64512
[R2-bgp]peer 10.0.12.1 as-number 64512

[R5]bgp 64514
[R5-bgp]undo peer 10.0.35.3
[R5-bgp]confederation id 200
[R5-bgp]confederation peer-as 64512
[R5-bgp]peer 10.0.35.3 as-number 64512
```

配置完成以后在R1和R3上分别检查对等体关系。

```
[R1]display bgp peer

 BGP local router ID : 10.0.1.1
 Local AS number : 64512
 Total number of peers : 3              Peers in established state : 3

  Peer            V          AS  MsgRcvd  MsgSent  OutQ  Up/Down       State PrefRcv

  10.0.12.2       4       64513        3        5     0 00:00:09 Established      1
  10.0.13.3       4       64512       22       24     0 00:18:16 Established      2
  10.0.14.4       4         100       25       27     0 00:22:39 Established      1

[R3]display bgp peer

 BGP local router ID : 10.0.13.3
 Local AS number : 64512
 Total number of peers : 2              Peers in established state : 2

  Peer            V          AS  MsgRcvd  MsgSent  OutQ  Up/Down       State PrefRcv

  10.0.13.1       4      64512       23       22     0 00:18:46 Established      3
  10.0.35.5       4      64514       10       12     0 00:07:32 Established      1
```

在R1上查看BGP路由表的详细信息，由于10.0.4.4/32这条路由是从EBGP对等体这里学习到的，所以标识为external。

```
[R1]display bgp routing-table 10.0.4.4

 BGP local router ID : 10.0.1.1
```

```
 Local AS number : 64512
 Paths:   1 available, 1 best, 1 select
 BGP routing table entry information of 10.0.4.4/32:
 From: 10.0.14.4 (10.0.14.4)
 Route Duration: 00h58m01s
 Direct Out-interface: Serial3/0/0
 Original nexthop: 10.0.14.4
 Qos information : 0x0
 AS-path 100, origin igp, MED 0, pref-val 0, valid, external, best, select, active,
pre 255
 Advertised to such 2 peers:
    10.0.13.3
    10.0.12.2
```

10.0.2.2/32这条路由是由R2始发的，R2与R1同属于AS 200，但属于联盟内不同的子AS，所以看到该路由的属性是external-confed，说明该路由从联盟内的EBGP对等体这里学习到的。

```
[R1]display bgp routing-table 10.0.2.2

 BGP local router ID : 10.0.1.1
 Local AS number : 64512
 Paths:   1 available, 1 best, 1 select
 BGP routing table entry information of 10.0.2.2/32:
 From: 10.0.12.2 (10.0.12.2)
 Route Duration: 00h36m05s
 Relay IP Nexthop: 0.0.0.0
 Relay IP Out-Interface: Serial1/0/0
 Original nexthop: 10.0.12.2
 Qos information : 0x0
 AS-path (64513), origin igp, MED 0, localpref 100, pref-val 0, valid,
external-confed, best, select, active, pre 255
 Advertised to such 2 peers:
    10.0.14.4
    10.0.13.3
```

最后观察路由10.0.3.3/32，该路由由R3始发，R3与R1同属于联盟内一个子AS，因此我们看到的该路由的属性是internal-confed，说明该路由从联盟内的IBGP对等体这里学习到的。

```
[R1]display bgp routing-table 10.0.3.3

 BGP local router ID : 10.0.1.1
 Local AS number : 64512
```

```
 Paths:   1 available, 1 best, 1 select
 BGP routing table entry information of 10.0.3.3/32:
 From: 10.0.13.3 (10.0.13.3)
 Route Duration: 00h53m23s
 Relay IP Nexthop: 0.0.0.0
 Relay IP Out-Interface: GigabitEthernet0/0/0
 Original nexthop: 10.0.13.3
 Qos information : 0x0
 AS-path Nil, origin igp, MED 0, localpref 100, pref-val 0, valid, internal-confed,
best, select, active, pre 255
 Advertised to such 2 peers:
    10.0.14.4
    10.0.12.2
```

步骤五. 观察联盟的属性

在部署了联盟以后，在子AS之间，BGP中的一些属性的特性会发生变化。我们首先观察AS-Path和Next-hop属性。

在R1上查看BGP路由表。

可以观察到对于从R4这里学习到的路由10.0.4.4/32，AS-Path属性的值为100。这些路由是从EBGP对等体那里学习到的路由。同时注意到路由10.0.2.2/32和10.0.5.5/32的AS-Path值带括号，这两条路由是R1从其他子AS学习到的。

我们可以通过AS-Path的值是否带括号来判断该路由是否是从其他子AS学习到的。

```
[R1]display bgp routing-table

 BGP Local router ID is 10.0.1.1
 Status codes: * - valid, > - best, d - damped,
               h - history,  i - internal, s - suppressed, S - Stale
               Origin : i - IGP, e - EGP, ? - incomplete


 Total Number of Routes: 5
      Network            NextHop        MED        LocPrf    PrefVal Path/Ogn

 *>   10.0.1.1/32         0.0.0.0         0                     0      i
 *>i  10.0.2.2/32         10.0.12.2       0          100        0      (64513)i
 *>i  10.0.3.3/32         10.0.13.3       0          100        0      i
 *>   10.0.4.4/32         10.0.14.4       0                     0      100i
```

```
  i  10.0.5.5/32        10.0.35.5       0          100        0      (64514)i
```

注意在R1上，10.0.5.5/32这条路由没有被BGP优选，即R1并不会将该路由通告给R4。

仔细观察，发现该路由的下一跳是10.0.35.5，该地址是R5连接到R3的互联地址。在R1的路由表中并没有到达该网段的路由。

检查R3的配置，已经在对等体10.0.13.1上配置了Next-hop-local，但我们可以发现，从R3发布过来的路由的下一跳并没有改变。

可以总结出，**对于联盟而言，子AS之间即使配置了Next-hop-local，也不改变路由的下一跳属性。**

从R3的路由表中也可以看到这一点，注意10.0.2.2/32和10.0.4.4/32这两条路由均是R3从其他对等体学习到的。

10.0.4.4/32是R1从AS 100学习到的，R1更改了下一跳为自己(10.0.13.1)，再将该路由发布给R3。

10.0.2.2/32是R1从R2这里学习到的，R2和R1同属于AS 200，但属于不同的子AS，所以R1并没有改变这条路由的下一跳，直接将该路由发布给了R3。

```
[R3]display bgp routing-table

 BGP Local router ID is 10.0.3.3
 Status codes: * - valid, > - best, d - damped,
               h - history,  i - internal, s - suppressed, S - Stale
               Origin : i - IGP, e - EGP, ? - incomplete


 Total Number of Routes: 5
      Network            NextHop        MED        LocPrf    PrefVal Path/Ogn

 *>i  10.0.1.1/32        10.0.13.1       0          100        0      i
   i  10.0.2.2/32        10.0.12.2       0          100        0      (64513)i
 *>   10.0.3.3/32        0.0.0.0         0                     0      i
 *>i  10.0.4.4/32        10.0.13.1       0          100        0      100i
 *>i  10.0.5.5/32        10.0.35.5       0          100        0      (64514)i
```

为了保证部署了联盟以后路由信息能够正常传递，把互联地址通过**network**发布出去

```
[R1-bgp]network 10.0.12.0 24
[R1-bgp]network 10.0.13.0 24

[R3-bgp]network 10.0.13.0 24
```

```
[R3-bgp]network 10.0.35.0 24
```

在R4上检查路由表，可以看到R4学习到了所有路由器的Loopback地址。同时还可以看到在AS-Path里，R4学习到的其他AS的路由的AS-Path均为200。即子AS号不会通告给联盟以外的AS，其他AS仅学习到联盟ID号。

```
[R4]display bgp routing-table

 BGP Local router ID is 10.0.14.4
 Status codes: * - valid, > - best, d - damped,
               h - history,  i - internal, s - suppressed, S - Stale
               Origin : i - IGP, e - EGP, ? - incomplete


 Total Number of Routes: 7
      Network            NextHop        MED        LocPrf    PrefVal Path/Ogn

 *>   10.0.1.1/32         10.0.14.1       0                     0      200i
 *>   10.0.2.2/32         10.0.14.1                             0      200i
 *>   10.0.3.3/32         10.0.14.1                             0      200i
 *>   10.0.4.4/32         0.0.0.0         0                     0      i
 *>   10.0.5.5/32         10.0.14.1                             0      200i
 *>   10.0.12.0/24        10.0.14.1       0                     0      200i
 *>   10.0.13.0/24        10.0.14.1       0                     0      200i
 *>   10.0.35.0/24        10.0.14.1                             0      200i
```

在R5上观察BGP路由表，可以看到从R4学习到的路由10.0.4.4/32是由AS 100始发的，然后经过子AS号64512，该AS号是带括号的。

这就说明AS 100为联盟以外的自治系统 ,AS 64512为联盟内的子自治系统。

```
[R5]display bgp routing-table

 BGP Local router ID is 10.0.5.5
 Status codes: * - valid, > - best, d - damped,
               h - history,  i - internal, s - suppressed, S - Stale
               Origin : i - IGP, e - EGP, ? - incomplete


 Total Number of Routes: 9
      Network            NextHop        MED        LocPrf    PrefVal Path/Ogn

 *>i  10.0.1.1/32         10.0.13.1       0          100        0      (64512)i
 *>i  10.0.2.2/32         10.0.12.2       0          100        0      (64512 64513)i
 *>i  10.0.3.3/32         10.0.35.3       0          100        0      (64512)i
```

```
*>i  10.0.4.4/32       10.0.13.1      0         100       0        (64512) 100i
*>   10.0.5.5/32       0.0.0.0         0                  0        i
*>i  10.0.12.0/24      10.0.13.1      0         100       0        (64512)i
*>i  10.0.13.0/24      10.0.35.3      0         100       0        (64512)i
*>i  10.0.35.0/24      10.0.35.3      0         100       0        (64512)i
```

在前面的实验中，我们已经介绍了团体属性中**no-advertise**和**no-export**两个特殊的值，其中**no-advertise**表示该路由不会通告给任何对等体，**no-export**表示该路由不能被发布到本地AS之外。

在联盟环境中，还有一种特殊的团体值：**no-export-subconfed**，该团体属性表示不能向AS外发送匹配的路由，也不能发布给其他子自治系统。

为了验证该特性，我们创建一个路由策略，将R5发布的路由10.0.5.5/32加上该团体属性，并允许在路由器之间传递团体属性，观察这条路由信息的传递过程。

首先在R5上创建访问控制列表2001，匹配R5发布的路由10.0.5.5/32。

```
[R5]acl number 2001
[R5-acl-basic-2001]rule permit source 10.0.5.5 0.0.0.0
[R5-acl-basic-2001]quit
```

接下来创建路由策略confed_community_control，针对匹配到ACL2001的路由，加上no-export-subconfed的属性。

```
[R5]route-policy confed_community_control permit node 10
Info: New Sequence of this List.
[R5-route-policy]if-match acl 2001
[R5-route-policy]apply community no-export-subconfed
[R5-route-policy]quit
```

然后将该路由策略应用到对等体10.0.35.3的出方向上，并允许向该对等体通告团体属性。然后重置R3与R5之间的对等体关系，确保路由策略已生效。

```
[R5]bgp 64514
[R5-bgp]peer 10.0.35.3 advertise-community
[R5-bgp]peer 10.0.35.3 route-policy confed_community_control export
[R5-bgp]return
<R5>refresh bgp all export
```

在R3上检查BGP路由表，发现它已学习到了该团体属性。在R3上将该团体属性发布给R1。

```
[R3]display bgp routing-table community
```

```
 BGP Local router ID is 10.0.3.3
 Status codes: * - valid, > - best, d - damped,
               h - history,  i - internal, s - suppressed, S - Stale
               Origin : i - IGP, e - EGP, ? - incomplete


 Total Number of Routes: 1
      Network          NextHop      MED      LocPrf    PrefVal Community

 *>i  10.0.5.5/32    10.0.35.5     0        100        0      no-export-subconfed
[R3]bgp 64512
[R3-bgp]peer 10.0.13.1 advertise-community
```

在R1上检查已学习到该团体属性以后，配置向其他对等体发送团体属性。

```
[R1]display bgp routing-table community

 BGP Local router ID is 10.0.1.1
 Status codes: * - valid, > - best, d - damped,
               h - history,  i - internal, s - suppressed, S - Stale
               Origin : i - IGP, e - EGP, ? - incomplete


 Total Number of Routes: 1
      Network            NextHop      MED      LocPrf    PrefVal Community

 *>i  10.0.5.5/32     10.0.35.5      0        100        0      no-export-subconfed
[R1]bgp 64512
[R1-bgp]peer 10.0.14.4 advertise-community
[R1-bgp]peer 10.0.12.2 advertise-community
```

在R2和R4上检查带团体属性的BGP路由表，可以发现均没有路由条目。

```
[R2]display bgp routing-table  community

 Total Number of Routes: 0

[R4]display bgp routing-table community

 Total Number of Routes: 0
```

这就证明了拥有no-export-subconfed团体属性的路由仅能在联盟内的子AS本地内部传递。

接下来我们再观察MED及本地优先属性在联盟内子AS的传递过程，在没有

联盟的环境中，这两个属性均是非传递的。

我们首先在R2上创建路由策略，给R2发布的10.0.2.2/32加上非默认的MED及本地优先级。

```
[R2]acl number 2001
[R2-acl-basic-2001]rule permit source 10.0.2.2 0.0.0.0
[R2-acl-basic-2001]quit
[R2]route-policy r2 permit node 10
Info: New Sequence of this List.
[R2-route-policy]if-match acl 2001
[R2-route-policy]apply cost 100
[R2-route-policy]apply local-preference 150
[R2-route-policy]quit
[R2]bgp 64513
[R2-bgp]peer 10.0.12.1 route-policy r2 export
```

重置R1与R2之间对等体关系之后，在R1上检查BGP路由表。

```
[R1]display bgp routing-table

 BGP Local router ID is 10.0.1.1
 Status codes: * - valid, > - best, d - damped,
               h - history,  i - internal, s - suppressed, S - Stale
               Origin : i - IGP, e - EGP, ? - incomplete


 Total Number of Routes: 8
      Network            NextHop        MED        LocPrf    PrefVal  Path/Ogn

 *>   10.0.1.1/32         0.0.0.0        0                      0      i
 *>i  10.0.2.2/32         10.0.12.2      100        150         0      (64513)i
 *>i  10.0.3.3/32         10.0.13.3      0          100         0      i
 *>   10.0.4.4/32         10.0.14.4      0                      0      100i
 *>i  10.0.5.5/32         10.0.35.5      0          100         0      (64514)i
 *>   10.0.12.0/24        0.0.0.0        0                      0      i
 *>   10.0.13.0/24        0.0.0.0        0                      0      i
   i                      10.0.13.3      0          100        0      i
 *>i  10.0.35.0/24        10.0.13.3      0          100         0      i
```

这时可以看到10.0.2.2/32这条路由的MED和本地优先级已经是在R2上设置的100和150。

分别在R4和R5上检查BGP路由表，观查这2个属性的传递情况。

```
[R4]display bgp routing-table
```

```
 BGP Local router ID is 10.0.4.4
 Status codes: * - valid, > - best, d - damped,
               h - history,  i - internal, s - suppressed, S - Stale
               Origin : i - IGP, e - EGP, ? - incomplete


 Total Number of Routes: 7
      Network            NextHop      MED        LocPrf    PrefVal Path/Ogn

 *>   10.0.1.1/32        10.0.14.1       0                     0      200i
 *>   10.0.2.2/32        10.0.14.1                             0      200i
 *>   10.0.3.3/32        10.0.14.1                             0      200i
 *>   10.0.4.4/32        0.0.0.0         0                     0      i
 *>   10.0.12.0/24       10.0.14.1       0                     0      200i
 *>   10.0.13.0/24       10.0.14.1       0                     0      200i
 *>   10.0.35.0/24       10.0.14.1                             0      200i

[R5]display bgp routing-table

 BGP Local router ID is 10.0.5.5
 Status codes: * - valid, > - best, d - damped,
               h - history,  i - internal, s - suppressed, S - Stale
               Origin : i - IGP, e - EGP, ? - incomplete


 Total Number of Routes: 8
      Network            NextHop     MED       LocPrf    PrefVal Path/Ogn

 *>i  10.0.1.1/32        10.0.13.1   0          100          0      (64512)i
 *>i  10.0.2.2/32        10.0.12.2   100        150          0      (64512 64513)i
 *>i  10.0.3.3/32        10.0.35.3   0          100          0      (64512)i
 *>i  10.0.4.4/32        10.0.13.1   0          100          0      (64512) 100i
 *>   10.0.5.5/32        0.0.0.0     0                       0      i
 *>i  10.0.12.0/24       10.0.13.1   0          100          0      (64512)i
 *>i  10.0.13.0/24       10.0.35.3   0          100          0      (64512)i
 *>i  10.0.35.0/24       10.0.35.3   0          100          0      (64512)i
```

从上面的输出中我们可以得知，对于MED和本地优先级，在联盟内部均可以正常传递，但不能被传递到联盟外部。

附加实验：思考并验证

在部署联盟的环境中，是否可以在子AS的边界上使用Aggregate对BGP路由进行汇总？

最终设备配置

```
<R1>display current-configuration
[V200R001C00SPC200]
#
 sysname R1
#
interface Serial1/0/0
 link-protocol ppp
 ip address 10.0.12.1 255.255.255.0
#
interface Serial3/0/0
 link-protocol ppp
 ip address 10.0.14.1 255.255.255.0
#
interface GigabitEthernet0/0/0
 ip address 10.0.13.1 255.255.255.0
#
interface LoopBack0
 ip address 10.0.1.1 255.255.255.255
#
bgp 64512
 router-id 10.0.1.1
 confederation id 200
 confederation peer-as 64513
 peer 10.0.12.2 as-number 64513
 peer 10.0.13.3 as-number 64512
 peer 10.0.14.4 as-number 100
 #
 ipv4-family unicast
  undo synchronization
  network 10.0.1.1 255.255.255.255
  network 10.0.12.0 255.255.255.0
  network 10.0.13.0 255.255.255.0
  peer 10.0.12.2 enable
```

```
  peer 10.0.12.2 advertise-community
  peer 10.0.13.3 enable
  peer 10.0.13.3 next-hop-local
  peer 10.0.13.3 advertise-community
  peer 10.0.14.4 enable
  peer 10.0.14.4 advertise-community
#
return

<R2>display current-configuration
[V200R001C00SPC200]
#
 sysname R2
#
acl number 2001
 rule 5 permit source 10.0.2.2 0
#
interface Serial1/0/0
 link-protocol ppp
 ip address 10.0.12.2 255.255.255.0
#
interface LoopBack0
 ip address 10.0.2.2 255.255.255.255
#
bgp 64513
 router-id 10.0.2.2
 confederation id 200
 confederation peer-as 64512
 peer 10.0.12.1 as-number 64512
 #
 ipv4-family unicast
  undo synchronization
  network 10.0.2.2 255.255.255.255
  peer 10.0.12.1 enable
  peer 10.0.12.1 route-policy r2 export
#
route-policy r2 permit node 10
 if-match acl 2001
 apply local-preference 150
 apply cost 100
#
return
```

```
<R3>display current-configuration
[V200R001C00SPC200]
#
 sysname R3
#
interface Serial3/0/0
 link-protocol ppp
 ip address 10.0.35.3 255.255.255.0
#
interface GigabitEthernet0/0/0
 ip address 10.0.13.3 255.255.255.0
#
interface LoopBack0
 ip address 10.0.3.3 255.255.255.255
#
bgp 64512
router-id 10.0.3.3
 confederation id 200
 confederation peer-as 64514
 peer 10.0.13.1 as-number 64512
 peer 10.0.35.5 as-number 64514
 #
 ipv4-family unicast
  undo synchronization
  network 10.0.3.3 255.255.255.255
  network 10.0.13.0 255.255.255.0
  network 10.0.35.0 255.255.255.0
  peer 10.0.13.1 enable
  peer 10.0.13.1 next-hop-local
  peer 10.0.13.1 advertise-community
  peer 10.0.35.5 enable
#
return
```

```
<R4>display current-configuration
[V200R001C00SPC200]
#
 sysname R4
#
interface Serial1/0/0
 link-protocol ppp
 ip address 10.0.14.4 255.255.255.0
#
```

```
interface LoopBack0
 ip address 10.0.4.4 255.255.255.255
#
bgp 100
router-id 10.0.4.4
 peer 10.0.14.1 as-number 200
 #
 ipv4-family unicast
  undo synchronization
  network 10.0.4.4 255.255.255.255
  peer 10.0.14.1 enable
#
return

<R5>display current-configuration
[V200R001C00SPC200]
#
 sysname R5
#
acl number 2001
 rule 5 permit source 10.0.5.5 0
#
interface Serial1/0/0
 link-protocol ppp
 ip address 10.0.35.5 255.255.255.0
#
interface LoopBack0
 ip address 10.0.5.5 255.255.255.255
#
bgp 64514
 router-id 10.0.5.5
 confederation id 200
 confederation peer-as 64512
 peer 10.0.35.3 as-number 64512
 #
 ipv4-family unicast
  undo synchronization
  network 10.0.5.5 255.255.255.255
  peer 10.0.35.3 enable
  peer 10.0.35.3 route-policy confed_community_control export
  peer 10.0.35.3 advertise-community
#
route-policy confed_community_control permit node 10
```

```
 if-match acl 2001
 apply community no-export-subconfed
#
return
```

第四章 访问控制

实验 4-1 应用 ACL 控制企业数据访问

学习目的

- 掌握基本ACL对数据进行过滤的方法
- 掌握调整基本ACL实现各种增强功能的方法
- 掌握高级ACL对数据进行过滤的方法
- 掌握调整高级ACL实现各种增强功能的方法

拓扑图

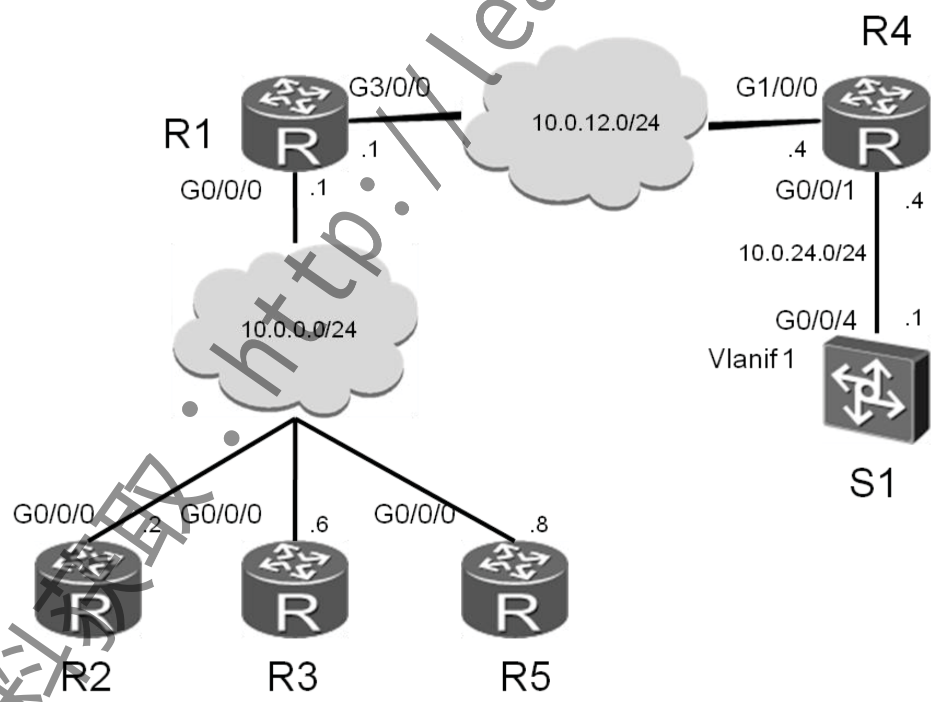

图4-1 应用ACL控制企业数据访问

场景

你是公司网络管理员。现在公司有两个网络，分别是总部和分部。总部的边界网络设备是R1，分部的边界网络设备是R4，它们之间通过Internet相连。公司需要控制员工的访问行为，需求是限制分部和总部的某些部门PC之间的互访。加强对Telnet访问的控制，只允许总部某些PC使用Telnet访问外网，同时不允许外网主动Telnet总部网络，以提高安全性。另外在总部的出口路由器上对内部员工访问外网的时间进行限制。R2、R3、R5和S1将模拟成客户端进行网络测试。

学习任务

步骤一. 基础配置与 IP 编址

给所有设备配置IP地址和掩码。

```
<Huawei>system-view
Enter system view, return user view with Ctrl+Z.
[Huawei]sysname R1
[R1]interface GigabitEthernet 0/0/0
[R1-GigabitEthernet0/0/0]ip address 10.0.0.1 24
[R1-GigabitEthernet0/0/0]interface Serial 3/0/0
[R1-Serial3/0/0]ip address 10.0.12.1 24

<Huawei>system-view
Enter system view, return user view with Ctrl+Z.
[Huawei]sysname R2
[R2]interface GigabitEthernet 0/0/0
[R2-GigabitEthernet0/0/0]ip address 10.0.0.2 24
[R2-GigabitEthernet0/0/0]quit
[R2]ip route-static 0.0.0.0 0 10.0.0.1

<Huawei>system-view
Enter system view, return user view with Ctrl+Z.
[Huawei]sysname R3
[R3]interface GigabitEthernet 0/0/0
[R3-GigabitEthernet0/0/0]ip address 10.0.0.6 24
[R3-GigabitEthernet0/0/0]quit
[R3]ip route-static 0.0.0.0 0 10.0.0.1

<Huawei>system-view
Enter system view, return user view with Ctrl+Z.
```

```
[Huawei]sysname R4
[R4]inter g0/0/1
[R4-GigabitEthernet0/0/1]ip address 10.0.24.4 24
[R4-GigabitEthernet0/0/1]interface Serial 1/0/0
[R4-Serial1/0/0]ip address 10.0.12.4 24

<Huawei>system-view
Enter system view, return user view with Ctrl+Z.
[Huawei]sysname R5
[R5]interface GigabitEthernet 0/0/0
[R5-GigabitEthernet0/0/0]ip address 10.0.0.8 24
[R5-GigabitEthernet0/0/0]quit
[R5]ip route-static 0.0.0.0 0 10.0.0.1

<Quidway>system-view
Enter system view, return user view with Ctrl+Z.
[Quidway]sysname S1
[S1]interface Vlanif 1
[S1-Vlanif1]ip address 10.0.24.1 24
[S1-Vlanif1]quit
[S1]ip route-static 0.0.0.0 0 10.0.24.4
```

在R2上测试与R3、R5、R1直连链路的连通性。

```
[R2]ping -c 1 10.0.0.6
  PING 10.0.0.6: 56  data bytes, press CTRL_C to break
    Reply from 10.0.0.6: bytes=56 Sequence=1 ttl=255 time=14 ms

  --- 10.0.0.6 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 14/14/14 ms

[R2]ping -c 1 10.0.0.8
  PING 10.0.0.8: 56  data bytes, press CTRL_C to break
    Reply from 10.0.0.8: bytes=56 Sequence=1 ttl=255 time=13 ms

  --- 10.0.0.8 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 13/13/13 ms
```

```
[R2]ping -c 1 10.0.0.1
  PING 10.0.0.1: 56  data bytes, press CTRL_C to break
    Reply from 10.0.0.1: bytes=56 Sequence=1 ttl=255 time=13 ms

  --- 10.0.0.1 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 13/13/13 ms
```

在R1上测试与R4直连链路的连通性。

```
[R1]ping -c 1 10.0.12.4
  PING 10.0.12.4: 56  data bytes, press CTRL_C to break
    Reply from 10.0.12.4: bytes=56 Sequence=1 ttl=255 time=42 ms

  --- 10.0.12.4 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 42/42/42 ms

    [R4]ping -c 1 10.0.24.1
      PING 10.0.24.1: 56 data bytes, press CTRL_C to break
        Reply from 10.0.24.1: bytes=56 Sequence=1 ttl=255 time=11 ms

      --- 10.0.24.1 ping statistics ---
        1 packet(s) transmitted
        1 packet(s) received
        0.00% packet loss
        round-trip min/avg/max = 11/11/11 ms
```

步骤二.　配置 OSPF 实现各网络互通

在R1和R4上配置OSPF协议，正确宣告它们各自连接的网络。此实验配置OSPF保证基本的网络互通，使用单区域的OSPF。

```
[R1]ospf 1 router-id 10.0.12.1
[R1-ospf-1]area 0
[R1-ospf-1-area-0.0.0.0]network 10.0.12.1 0.0.0.0
[R1-ospf-1-area-0.0.0.0]network 10.0.0.1 0.0.0.0

[R4]ospf 1 router-id 10.0.12.4
```

```
[R4-ospf-1]area 0
[R4-ospf-1-area-0.0.0.0]network 10.0.12.4 0.0.0.0
[R4-ospf-1-area-0.0.0.0]network 10.0.24.4 0.0.0.0
```

查看R1和R4的路由表。

```
[R1]disp ip routing-table
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Routing Tables: Public
        Destinations : 12      Routes : 12

Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface

      10.0.0.0/24  Direct  0    0           D   10.0.0.1        GigabitEthernet0/0/0
      10.0.0.1/32  Direct  0    0           D   127.0.0.1       InLoopBack0
    10.0.0.255/32  Direct  0    0           D   127.0.0.1       InLoopBack0
     10.0.12.0/24  Direct  0    0           D   10.0.12.1       Serial3/0/0
     10.0.12.1/32  Direct  0    0           D   127.0.0.1       InLoopBack0
     10.0.12.4/32  Direct  0    0           D   10.0.12.4       Serial3/0/0
   10.0.12.255/32  Direct  0    0           D   127.0.0.1       InLoopBack0
     10.0.24.0/24  OSPF    10   1563        D   10.0.12.4       Serial3/0/0
     127.0.0.0/8   Direct  0    0           D   127.0.0.1       InLoopBack0
     127.0.0.1/32  Direct  0    0           D   127.0.0.1       InLoopBack0
127.255.255.255/32  Direct  0    0           D   127.0.0.1       InLoopBack0
255.255.255.255/32  Direct  0    0           D   127.0.0.1       InLoopBack0


[R4]disp ip routing-table
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Routing Tables: Public
        Destinations : 12      Routes : 12

Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface

      10.0.0.0/24  OSPF    10   1563        D   10.0.12.1       Serial1/0/0
     10.0.12.0/24  Direct  0    0           D   10.0.12.4       Serial1/0/0
     10.0.12.1/32  Direct  0    0           D   10.0.12.1       Serial1/0/0
     10.0.12.4/32  Direct  0    0           D   127.0.0.1       InLoopBack0
   10.0.12.255/32  Direct  0    0           D   127.0.0.1       InLoopBack0
     10.0.24.0/24  Direct  0    0           D   10.0.24.4       GigabitEthernet0/0/1
     10.0.24.4/32  Direct  0    0           D   127.0.0.1       InLoopBack0
   10.0.24.255/32  Direct  0    0           D   127.0.0.1       InLoopBack0
```

```
     127.0.0.0/8   Direct  0    0          D   127.0.0.1       InLoopBack0
     127.0.0.1/32  Direct  0    0          D   127.0.0.1       InLoopBack0
127.255.255.255/32  Direct  0    0          D   127.0.0.1       InLoopBack0
255.255.255.255/32  Direct  0    0          D   127.0.0.1       InLoopBack0
```

从R1和R4输出的路由表可以看到已经正确学习到对方的网络信息。

在R2使用**ping**命令测试与S1的连通性。

```
[R2]ping -c 1 10.0.24.1
  PING 10.0.24.1: 56  data bytes, press CTRL_C to break
    Reply from 10.0.24.1: bytes=56 Sequence=1 ttl=253 time=31 ms

  --- 10.0.24.1 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 31/31/31 ms
```

在S1上使用**ping**命令测试与R3的连通性。

```
[S1]ping -c 1 10.0.0.6
     PING 10.0.0.6: 56  data bytes, press CTRL_C to break
       Reply from 10.0.0.6: bytes=56 Sequence=1 ttl=253 time=30 ms

     --- 10.0.0.6 ping statistics ---
       1 packet(s) transmitted
       1 packet(s) received
       0.00% packet loss
       round-trip min/avg/max = 30/30/30 ms
```

步骤三.　基本 ACL 的配置

限制分部的S1与总部的R2、R3、R5之间的互访。在这个步骤使用基本ACL实现。

在配置ACL时，可以通过Match-order参数来控制ACL语句的顺序，以便实现最大的配置的可扩展性。

在R4上配置基本ACL阻止R2、R3、R5访问S1。

```
[R4]acl 2000
[R4-acl-basic-2000]rule deny source 10.0.0.0 0.0.0.255
[R4-acl-basic-2000]rule permit source any
```

使用**display acl all**命令检查ACL配置结果。

```
[R4]disp acl all
 Total nonempty ACL number is 1

Basic ACL 2000, 2 rules
Acl's step is 5
 rule 5 deny source 10.0.0.0 0.0.0.255
 rule 10 permit
```

为R4建立Outside区域和Inside区域，并将优先级分别设置为1和10。

```
[R4]firewall zone outside
[R4-zone-outside]priority 1
[R4-zone-outside]quit
[R4]firewall zone inside
[R4-zone-inside]priority 10
```

将R4的S1/0/0接口加入Outside区域，G0/0/1接口加入Inside区域。

```
[R4]interface Serial 1/0/0
[R4-Serial1/0/0]zone outside
[R4-Serial1/0/0]inter g0/0/1
[R4-GigabitEthernet0/0/1]zone inside
```

在R4上配置基于ACL的包过滤。

```
[R4]firewall interzone inside outside
[R4-interzone-inside-outside]packet-filter 2000 inbound
[R4-interzone-inside-outside]firewall enable
```

使用**display firewall interzone**命令检查配置结果。

```
[R4]display firewall interzone inside outside
interzone inside outside
 firewall enable
 packet-filter default deny inbound
 packet-filter default permit outbound
 packet-filter 2000 inbound
```

在R2、R3、R5上测试与R4的连通性。

```
[R2]ping -c 1 10.0.24.1
  PING 10.0.24.1: 56  data bytes, press CTRL_C to break
    Request time out
```

```
  --- 10.0.24.1 ping statistics ---
    1 packet(s) transmitted
    0 packet(s) received
100.00% packet loss

[R3]ping -c 1 10.0.24.1
  PING 10.0.24.1: 56  data bytes, press CTRL_C to break
    Request time out

  --- 10.0.24.1 ping statistics ---
    1 packet(s) transmitted
    0 packet(s) received
    100.00% packet loss
[R5]ping -c 1 10.0.24.1
  PING 10.0.24.1: 56 data bytes, press CTRL_C to break
    Request time out

  --- 10.0.24.1 ping statistics ---
    1 packet(s) transmitted
    0 packet(s) received
    100.00% packet loss

[R1]ping -c 1 10.0.24.1
  PING 10.0.24.1: 56  data bytes, press CTRL_C to break
    Reply from 10.0.24.1: bytes=56 Sequence=1 ttl=254 time=35 ms

  --- 10.0.24.1 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 35/35/35 ms
```

可见，基本ACL已经生效。R2、R3、R5均不能再与S1通信。而不受控制的R1仍然可以和S1通信。

在不改变原ACL的情况下添加新的ACL条目 ,实现在R5不能访问R4的前提下，R2、R3可以访问R4的结果。

在R4的ACL2000上添加允许R2、R3访问R4的ACL规则。

```
[R4]acl 2000
[R4-acl-basic-2000]rule permit source 10.0.0.2 0.0.0.252
```

使用**display acl all**命令检查ACL配置结果。

```
[R4-acl-basic-2000]disp acl all
 Total nonempty ACL number is 1

Basic ACL 2000, 3 rules
Acl's step is 5
 rule 5 deny source 10.0.0.0 0.0.0.255
 rule 10 permit
 rule 15 permit source 10.0.0.2 0.0.0.252
```

使用Match-order改变ACL的执行顺序。

```
[R4]acl 2000 match-order config
[R4-acl-basic-2000]undo rule 15
[R4-acl-basic-2000]rule 1 permit source 10.0.0.2 0.0.0.252
```

使用**display acl all**命令检查ACL配置结果。

```
[R4]disp acl all
 Total nonempty ACL number is 1

Basic ACL 2000, 3 rules
Acl's step is 5
 rule 1 permit source 10.0.0.2 0.0.0.252
 rule 5 deny source 10.0.0.0 0.0.0.255
 rule 10 permit
```

从输出的信息可以看到原来Rule 15的条目已经变为Rule 1。

在R2、R3、R5上测试与S1的连通性。

```
[R2]ping -c 1 10.0.24.1
  PING 10.0.24.1: 56  data bytes, press CTRL_C to break
    Reply from 10.0.24.1: bytes=56 Sequence=1 ttl=253 time=34 ms

  --- 10.0.24.1 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
round-trip min/avg/max = 34/34/34 ms

[R3]ping -c 1 10.0.24.1
  PING 10.0.24.1: 56  data bytes, press CTRL_C to break
    Reply from 10.0.24.1: bytes=56 Sequence=1 ttl=253 time=37 ms

  --- 10.0.24.1 ping statistics ---
```

```
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 37/37/37 ms

[R5]ping -c 1 10.0.24.1
  PING 10.0.24.1: 56 data bytes, press CTRL_C to break
    Request time out

  --- 10.0.24.1 ping statistics ---
    1 packet(s) transmitted
    0 packet(s) received
    100.00% packet loss
```

步骤四. 高级 ACL 的配置

可以使用高级ACL对Telnet应用做更加严格的访问控制。在配置高级ACL时，通过使用对TCP数据包的SYN-flag位的检查来实现。

在R4上启用Telnet服务。

```
[R4]user-interface vty 0 4
[R4-ui-vty0-4]authentication-mode password
[R4-ui-vty0-4]set authentication password simple huawei
```

在R2、R3上测试R4的Telnet服务。

```
<R2>telnet 10.0.12.4
  Press CTRL_] to quit telnet mode
  Trying 10.0.12.4 ...
  Connected to 10.0.12.4 ...

Login authentication


Password:
<R4>quit
  Configuration console exit, please retry to log on

  The connection was closed by the remote host
<R2>

<R3>telnet 10.0.12.4
  Press CTRL_] to quit telnet mode
```

```
  Trying 10.0.12.4 ...
  Connected to 10.0.12.4 ...

Login authentication


Password:
<R4>quit
  Configuration console exit, please retry to log on

  The connection was closed by the remote host
<R3>
```

在R1上配置高级ACL，只允许R2使用Telnet登陆R4的S1/0/0接口。

```
[R1]acl 3000
[R1-acl-adv-3000]rule  permit tcp source 10.0.0.2 0 destination-port eq telnet
[R1-acl-adv-3000]rule deny tcp destination-port eq telnet
```

检查ACL配置结果。

```
[R1]display acl all
 Total nonempty ACL number is 1

Advanced ACL 3000, 2 rules
Acl's step is 5
 rule 5 permit tcp source 10.0.0.2 0 destination-port eq telnet
 rule 10 deny tcp destination-port eq telnet
```

在R1上配置基于ACL的包过滤。

```
[R1]firewall zone out
[R1-zone-out]priority 1
[R1-zone-out]quit
[R1]firewall zone int
[R1-zone-int]priority 10
[R1-zone-int]quit
[R1]inter s3/0/0
[R1-Serial3/0/0]zone out
[R1-Serial3/0/0]inter g0/0/0
[R1-GigabitEthernet0/0/0]zone int
[R1-GigabitEthernet0/0/0]quit
[R1]firewall interzone int out
[R1-interzone-int-out]packet-filter 3000 outbound
[R1-interzone-int-out]firewall enable
```

使用**display firewall interzone**命令检查配置结果。

```
[R1]display firewall interzone int out
interzone int out
 firewall enable
 packet-filter default deny inbound
 packet-filter default permit outbound
 packet-filter 3000 outbound
```

测试R2、R3能否访问R4的S1/0/0接口提供的Telnet服务。

```
<R2>telnet 10.0.12.4
  Press CTRL_] to quit telnet mode
  Trying 10.0.12.4 ...
  Connected to 10.0.12.4 ...

Login authentication



Password:
<R4>quit
  Configuration console exit, please retry to log on
  The connection was closed by the remote host
<R2>

<R3>telnet 10.0.12.4
  Press CTRL_] to quit telnet mode
  Trying 10.0.12.4 ...
  Error: Can't connect to the remote host
```

在R2上启用Telnet服务。

```
[R2]user-interface vty 0 4
[R2-ui-vty0-4]authentication-mode password
[R2-ui-vty0-4]set authentication password simple huawei
```

因为R1的防火墙功能阻止了外网访问内容的主动连接，所以S1不能和R2通信。

```
<S1>ping -c 1 10.0.0.2
  PING 10.0.0.2: 56  data bytes, press CTRL_C to break
    Request time out

  --- 10.0.0.2 ping statistics ---
```

```
    1 packet(s) transmitted
    0 packet(s) received
    100.00% packet loss
```

调整R1上的防火墙策略，允许外网对企业内网的访问。

```
[R1]firewall interzone int out
[R1-interzone-int-out]packet-filter default permit inbound
```

在S1上测试R2的Telnet服务。

```
<S1>telnet 10.0.0.2
Trying 10.0.0.2 ...
Press CTRL+K to abort
Connected to 10.0.0.2 ...

Login authentication



Password:
<R2>quit
  Configuration console exit, please retry to log on

Info: The connection was closed by the remote host.
<S1>
```

在R1上配置高级ACL，实现只允许总部内网用户Telnet访问外网，不允许外网用户主动Telnet访问内网。

```
[R1]acl 3100
[R1-acl-adv-3100]rule deny tcp destination-port eq 23 tcp-flag ack
```

检查ACL配置结果。

```
[R1]display acl all
 Total nonempty ACL number is 2

Advanced ACL 3000, 2 rules
Acl's step is 5
 rule 5 permit tcp source 10.0.0.2 0 destination-port eq telnet
 rule 10 deny tcp destination-port eq telnet

Advanced ACL 3100, 1 rule
Acl's step is 5
 rule 5 deny tcp destination-port eq telnet tcp-flag ack
```

在R1上配置基于ACL的包过滤。

```
[R1]firewall interzone int out
[R1-interzone-int-out]packet-filter 3100 inbound
```

使用**display firewall interzone**命令检查配置结果。

```
[R1]disp fire inte int out
interzone int out
 firewall enable
 packet-filter default permit outbound
 packet-filter 3100 inbound
 packet-filter 3000 outbound
 packet-filter default permit inbound
```

分别测试在R2通过Telnet访问R4。S1通过Telnet访问R2。

```
<R2>telnet 10.0.12.4
Trying 10.0.12.4 ...
Press CTRL+K to abort
Connected to 10.0.12.4 ...

Login authentication


Password:
<R4>quit
  Configuration console exit, please retry to log on

Info: The connection was closed by the remote host.
<R2>

<S1>telnet 10.0.0.2
Trying 10.0.0.2 ...
Press CTRL+K to abort
Error: Failed to connect to the remote host.
```

在R1上使用**Debug**命令调试S1使用Telnet访问R2的过程。

```
<S1>terminal debugging
<S1>terminal monitor
<S1>debugging tcp packet

<S1>telnet 10.0.0.2
Trying 10.0.0.2 ...
```

```
Press CTRL+K to abort
Dec 15 2011 16:36:25.370.1-05:13 S1 SOCKET/7/TCP PACKET:
TCP debug packet information:
1323966985: Output: task = co0 (6), socketid = 1,
(State:Syn_Sent,src = 10.0.24.1:51451,dst = 10.0.0.2:23,VrfIndex = 0,seq =
3774177805,
ack = 0,datalen = 0,optlen = 4,flag =  SYN ,window = 8192,ttl = 255,tos = 0,MSS
= 512)

Dec 15 2011 16:36:31.40.1-05:13 S1 SOCKET/7/TCP PACKET:
TCP debug packet information:
1323966991: Output: task = co0 (6), socketid = 1,
(State:Syn_Sent,src = 10.0.24.1:51451,dst = 10.0.0.2:23,VrfIndex = 0,seq =
3774177805,
ack = 0,datalen = 0,optlen = 4,flag =  SYN ,window = 8192,ttl = 255,tos = 0,MSS
= 512)

Dec 15 2011 16:36:55.40.1-05:13 S1 SOCKET/7/TCP PACKET:
TCP debug packet information:
1323967015: Output: task = co0 (6), socketid = 1,
(State:Syn_Sent,src = 10.0.24.1:51451,dst = 10.0.0.2:23,VrfIndex = 0,seq =
3774177805,
ack = 0,datalen = 0,optlen = 4,flag =  SYN ,window = 8192,ttl = 255,tos = 0,MSS
= 512)

Error: Failed to connect to the remote host.
```

从**debug**命令输出结果可以看到，只有TCP连接的SYN数据包，没有ACK等其他类型的数据包。这是因为在R1上的ACL3100拒绝回应任何来自外部网络的TCP连接。

步骤五. 带有时间段 ACL 的配置

使用带有时间段的ACL可以实现基于时间的访问控制。

在R1上配置带有时间段的ACL，实现允许内网R2在工作日的8:00-18:00和周六23:00-周日的1:00时间内使用Telnet访问外网。

创建时间段。

```
[R1]time-range Telnet_Control 8:00 to 18:00 working-day
[R1]time-range Telnet_Control 23:00 to 0:00 Sat
[R1]time-range Telnet_Control 0:00 to 1:00 Sun
```

使用**disp time-range all**命令查看设置的时间段。

```
[R1]disp time-range all
Current time is 17:31:30 12-15-2011 Thursday

Time-range : Telnet_Control ( Active )
 08:00 to 18:00 working-day
 23:00 to 00:00 Sat
 00:00 to 01:00 Sun
```

在R1上配置允许内网对外网的访问并引入时间段。

```
[R1]acl 3000
[R1-acl-adv-3000]rule 5 permit tcp source 10.0.0.2 0 destination-port eq telnet
time-range Telnet_Control
```

检查ACL配置结果。

```
[R1-acl-adv-3000]disp acl 3000
Advanced ACL 3000, 2 rules
Acl's step is 5
 rule 5 permit tcp source 10.0.0.2 0 destination-port eq telnet time-range
Telnet_Control (Active)
 rule 10 deny tcp destination-port eq telnet
```

查看当前时间下R2使用Telnet访问外网的连通性。（实验时根据自己的实际时间，只要能够验证处于Telnet_Control规定的时间段时能够使用Telnet登陆R4，不处于规定时间段时不能使用Telnet登陆R4即可。）

```
<R1>display clock
2011-12-15 17:44:16
Thursday
Time Zone(Default Zone Name) : UTC+00:00
```

测试当前配置下R2使用Telnet访问外网的连通性。

```
<R2>telnet 10.0.12.4
  Press CTRL_] to quit telnet mode
  Trying 10.0.12.4 ...
  Connected to 10.0.12.4 ...

Login authentication


Password:
<R4>quit
```

```
  Configuration console exit, please retry to log on

  The connection was closed by the remote host
<R2>
```

修改系统时间到允许时间范围之外。

```
<R1>clock datetime 10:0:0 2011-12-17
<R1>disp clock
2011-12-17 10:00:05
```

测试当前配置下R2使用Telnet访问外网的连通性。

```
<R2>telnet 10.0.12.4
  Press CTRL_] to quit telnet mode
  Trying 10.0.12.4 ...
  Error: Can't connect to the remote host
```

修改系统时间到周六允许时间之内。

```
<R1>clock datetime 23:10:0 2011-12-17
<R1>disp clock
2011-12-17 23:10:02
```

测试当前配置下R2使用Telnet访问外网的连通性。

```
<R2>telnet 10.0.12.4
  Press CTRL_] to quit telnet mode
  Trying 10.0.12.4 ...
  Connected to 10.0.12.4 ...

Login authentication


Password:
<R4>quit
  Configuration console exit, please retry to log on

  The connection was closed by the remote host
<R2>
```

附加实验：思考并验证

请从网络优化角度思考在内网中应用基本ACL和高级ACL的时候，基本ACL的放置位置应该靠近源网络还是靠近目的网络，高级ACL的放置位置应该靠近源网络还是靠近目的网络，为什么？

最终设备配置

```
[R1]display current-configuration
[V200R001C00SPC200]
#
 sysname R1
#
 time-range telnet_control 08:00 to 18:00 working-day
 time-range telnet_control 23:00 to 00:00 Sat
 time-range telnet_control 00:00 to 01:00 Sun
#
acl number 3000
 rule 5 permit tcp source 10.0.0.2 0 destination-port eq telnet time-range
telnet_control
 rule 10 deny tcp destination-port eq telnet
#
acl number 3100
 rule 5 deny tcp destination-port eq telnet tcp-flag ack
#
firewall zone int
 priority 10
#
firewall zone out
 priority 1
#
firewall interzone int out
 firewall enable
 packet-filter 3100 inbound
 packet-filter 3000 outbound
 packet-filter default permit inbound
#
interface Serial3/0/0
 link-protocol ppp
 ip address 10.0.12.1 255.255.255.0
```

```
 zone out
#
interface GigabitEthernet0/0/0
 ip address 10.0.0.1 255.255.255.0
 zone int
#
ospf 1 router-id 10.0.12.1
 area 0.0.0.0
  network 10.0.12.1 0.0.0.0
  network 10.0.0.1 0.0.0.0
#
Return

[R2]display current-configuration
[V200R001C00SPC500]
#
 sysname R2
#
interface GigabitEthernet0/0/0
 ip address 10.0.0.2 255.255.255.0
#
ip route-static 0.0.0.0 0 10.0.0.1
#
return

[R3]display current-configuration
[V200R001C00SPC500]
#
 sysname R3
#
interface GigabitEthernet0/0/0
 ip address 10.0.0.6 255.255.255.0
#
ip route-static 0.0.0.0 0 10.0.0.1
#
return

[R4]display current-configuration
[V200R001C00SPC500]
#
 sysname R4
#
acl number 2000
```

```
 rule 1 permit source 10.0.0.2 0.0.0.252
 rule 5 deny source 10.0.0.0 0.0.0.255
 rule 10 permit
#
firewall zone inside
 priority 10
#
firewall zone outside
 priority 1
#
firewall interzone inside outside
 firewall enable
 packet-filter 2000 inbound
#
interface Serial1/0/0
 link-protocol ppp
 ip address 10.0.12.4 255.255.255.0
 zone outside
#
interface GigabitEthernet0/0/0
#
interface GigabitEthernet0/0/1
 ip address 10.0.24.4 255.255.255.0
 zone inside
#
ospf 1 router-id 10.0.12.4
 area 0.0.0.0
  network 10.0.12.4 0.0.0.0
  network 10.0.24.4 0.0.0.0
#
user-interface vty 0 4
 set authentication password simple huawei
user-interface vty 16 20
#
Return

[R5]display current-configuration
[V200R001C00SPC500]
#
 sysname R5
#
interface GigabitEthernet0/0/0
 ip address 10.0.0.8 255.255.255.0
```

```
#
ip route-static 0.0.0.0 0 10.0.0.1
#
return


[S1]display current-configuration
#
!Software Version V100R005C01SPC100
 sysname S1
#
interface Vlanif1
 ip address 10.0.24.1 255.255.255.0
#
ip route-static 0.0.0.0 0 10.0.24.4
#
Return
```

实验 4-2 路由引入与路由控制

学习目的

- 掌握OSPF与RIP相互路由引入的配置方法
- 掌握通过地址前缀列表过滤路由信息的配置方法
- 掌握通过Route-policy过滤路由信息的配置方法

拓扑图

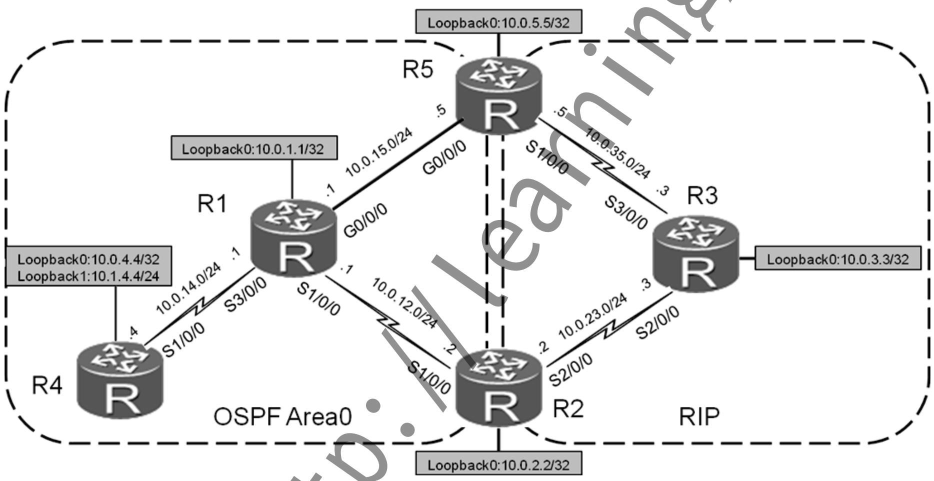

图4-2 路由引入与路由控制

场景

你是你们公司的网络管理员。公司网络中有两部分路由区域，一部分运行OSPF,另外一部分运行RIP。为了实现网络互通，你需要配置路由相互引入。在典型的双点双向路由引入中，会出现一些问题。为了解决这些问题，避免可能出现的路由环路和次优路径的产生，现在你需要使用到前缀列表和路由策略对路由进行控制。

学习任务

步骤一. 基础配置与 IP 编址

给所有路由器配置物理接口以及Loopback接口的IP地址和掩码。注意各Loopback 0接口均使用32位掩码。

```
<R1>system-view
Enter system view, return user view with Ctrl+Z.
[R1]interface  Serial 1/0/0
[R1-Serial1/0/0]ip address 10.0.12.1 255.255.255.0
[R1-Serial1/0/0]interface GigabitEthernet 0/0/0
[R1-GigabitEthernet0/0/0]ip address 10.0.15.1 255.255.255.0
[R1-GigabitEthernet0/0/0]interface Serial 3/0/0
[R1-Serial3/0/0]ip add 10.0.14.1 255.255.255.0
[R1-Serial3/0/0]int LoopBack 0
[R1-LoopBack0]ip address 10.0.1.1 255.255.255.255

<R2>system-view
Enter system view, return user view with Ctrl+Z.
[R2]interface Serial 1/0/0
[R2-Serial1/0/0]ip address 10.0.12.2 255.255.255.0
[R2-Serial1/0/0]interface Serial 2/0/0
[R2-Serial2/0/0]ip address 10.0.23.2 255.255.255.0
[R2-Serial2/0/0]interface LoopBack 0
[R2-LoopBack0]ip add 10.0.2.2 255.255.255.255

<R3>system-view
Enter system view, return user view with Ctrl+Z.
[R3]interface Serial 2/0/0
[R3-Serial2/0/0]ip address 10.0.23.3 255.255.255.0
[R3-Serial2/0/0]int Serial 3/0/0
[R3-Serial3/0/0]ip address 10.0.35.3 255.255.255.0
[R3-Serial3/0/0]interface LoopBack 0
[R3-LoopBack0]ip address 10.0.3.3 255.255.255.255

<R4>system-view
Enter system view, return user view with Ctrl+Z.
[R4]int Serial 1/0/0
[R4-Serial1/0/0]ip address 10.0.14.4 255.255.255.0
[R4-Serial1/0/0]int LoopBack 0
[R4-LoopBack0]ip add 10.0.4.4 255.255.255.255
```

```
<R5>system-view
Enter system view, return user view with Ctrl+Z.
[R5]int Serial 1/0/0
[R5-Serial1/0/0]ip address 10.0.35.5 255.255.255.0
[R5-Serial1/0/0]int GigabitEthernet 0/0/0
[R5-GigabitEthernet0/0/0]ip address 10.0.15.5 255.255.255.0
[R5-GigabitEthernet0/0/0]int LoopBack 0
[R5-LoopBack0]ip address 10.0.5.5 255.255.255.255
```

配置完成后，测试直连链路的连通性。

```
[R1]ping -c 1 10.0.12.2
  PING 10.0.12.2: 56  data bytes, press CTRL_C to break
    Reply from 10.0.12.2: bytes=56 Sequence=1 ttl=255 time=38 ms

  --- 10.0.12.2 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 38/38/38 ms

[R1]ping -c 1 10.0.15.5
  PING 10.0.15.5: 56  data bytes, press CTRL_C to break
    Reply from 10.0.15.5: bytes=56 Sequence=1 ttl=255 time=12 ms

  --- 10.0.15.5 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 12/12/12 ms

[R1]ping -c 1 10.0.14.4
  PING 10.0.14.4: 56  data bytes, press CTRL_C to break
    Reply from 10.0.14.4: bytes=56 Sequence=1 ttl=255 time=33 ms

  --- 10.0.14.4 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 33/33/33 ms

[R2]ping -c 1 10.0.23.3
  PING 10.0.23.3: 56  data bytes, press CTRL_C to break
```

```
    Reply from 10.0.23.3: bytes=56 Sequence=1 ttl=255 time=34 ms

  --- 10.0.23.3 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 34/34/34 ms

[R3]ping -c 1 10.0.35.5
  PING 10.0.35.5: 56  data bytes, press CTRL_C to break
    Reply from 10.0.35.5: bytes=56 Sequence=1 ttl=255 time=39 ms

  --- 10.0.35.5 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 39/39/39 ms
```

步骤二. 配置 IGP

在R1，R2，R4，R5上运行OSPF，所有设备属于区域0。

配置R1的S1/0/0，S3/0/0，G0/0/0和Loopback0连接的网段运行OSPF，

```
[R1]ospf 1
[R1-ospf-1]area 0
[R1-ospf-1-area-0.0.0.0]network 10.0.12.1 0.0.0.0
[R1-ospf-1-area-0.0.0.0]network 10.0.15.1 0.0.0.0
[R1-ospf-1-area-0.0.0.0]network 10.0.14.1 0.0.0.0
[R1-ospf-1-area-0.0.0.0]network 10.0.1.1 0.0.0.0
```

配置R2的S1/0/0连接的网段运行OSPF。

```
[R2]ospf 1
[R2-ospf-1]area 0
[R2-ospf-1-area-0.0.0.0]network 10.0.12.2 0.0.0.0
```

配置R4的S1/0/0和Loopback 0连接的网段运行OSPF。

```
[R4]ospf 1
[R4-ospf-1]area 0
[R4-ospf-1-area-0.0.0.0]network 10.0.14.4 0.0.0.0
[R4-ospf-1-area-0.0.0.0]network 10.0.4.4 0.0.0.0
```

配置R5的G0/0/0连接的网段运行OSPF。

```
[R5]ospf 1
[R5-ospf-1]area 0
[R5-ospf-1-area-0.0.0.0]network 10.0.15.5 0.0.0.0
```

检查是否学习到其他设备上的Loopback0接口连接网段的路由。

```
[R1]display ip routing-table
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Routing Tables: Public
        Destinations : 17       Routes : 17

Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface

      10.0.1.1/32   Direct  0    0           D   127.0.0.1       InLoopBack0
      10.0.4.4/32   OSPF    10   1562        D   10.0.14.4       Serial3/0/0
     10.0.12.0/24   Direct  0    0           D   10.0.12.1       Serial1/0/0
     10.0.12.1/32   Direct  0    0           D   127.0.0.1       InLoopBack0
     10.0.12.2/32   Direct  0    0           D   10.0.12.2       Serial1/0/0
   10.0.12.255/32   Direct  0    0           D   127.0.0.1       InLoopBack0
     10.0.14.0/24   Direct  0    0           D   10.0.14.1       Serial3/0/0
     10.0.14.1/32   Direct  0    0           D   127.0.0.1       InLoopBack0
     10.0.14.4/32   Direct  0    0           D   10.0.14.4       Serial3/0/0
   10.0.14.255/32   Direct  0    0           D   127.0.0.1       InLoopBack0
     10.0.15.0/24   Direct  0    0           D   10.0.15.1   GigabitEthernet0/0/0
     10.0.15.1/32   Direct  0    0           D   127.0.0.1       InLoopBack0
   10.0.15.255/32   Direct  0    0           D   127.0.0.1       InLoopBack0
     127.0.0.0/8    Direct  0    0           D   127.0.0.1       InLoopBack0
     127.0.0.1/32   Direct  0    0           D   127.0.0.1       InLoopBack0
127.255.255.255/32  Direct  0    0           D   127.0.0.1       InLoopBack0
255.255.255.255/32  Direct  0    0           D   127.0.0.1       InLoopBack0

[R2]display ip routing-table
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Routing Tables: Public
        Destinations : 17       Routes : 17

Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface

      10.0.1.1/32   OSPF    10   1562        D   10.0.12.1       Serial1/0/0
      10.0.2.2/32   Direct  0    0           D   127.0.0.1       InLoopBack0
```

```
      10.0.4.4/32   OSPF    10   3124        D   10.0.12.1      Serial1/0/0
     10.0.12.0/24   Direct  0    0           D   10.0.12.2      Serial1/0/0
     10.0.12.1/32   Direct  0    0           D   10.0.12.1      Serial1/0/0
     10.0.12.2/32   Direct  0    0           D   127.0.0.1      InLoopBack0
   10.0.12.255/32   Direct  0    0           D   127.0.0.1      InLoopBack0
     10.0.14.0/24   OSPF    10   3124        D   10.0.12.1      Serial1/0/0
     10.0.15.0/24   OSPF    10   1563        D   10.0.12.1      Serial1/0/0
     10.0.23.0/24   Direct  0    0           D   10.0.23.2      Serial2/0/0
     10.0.23.2/32   Direct  0    0           D   127.0.0.1      InLoopBack0
     10.0.23.3/32   Direct  0    0           D   10.0.23.3      Serial2/0/0
   10.0.23.255/32   Direct  0    0           D   127.0.0.1      InLoopBack0
      127.0.0.0/8   Direct  0    0           D   127.0.0.1      InLoopBack0
     127.0.0.1/32   Direct  0    0           D   127.0.0.1      InLoopBack0
127.255.255.255/32  Direct  0    0           D   127.0.0.1      InLoopBack0
255.255.255.255/32  Direct  0    0           D   127.0.0.1      InLoopBack0

[R4]display ip routing-table
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Routing Tables: Public
         Destinations : 12      Routes : 12

Destination/Mask    Proto   Pre  Cost       Flags NextHop         Interface

      10.0.1.1/32   OSPF    10   1562        D   10.0.14.1      Serial1/0/0
      10.0.4.4/32   Direct  0    0           D   127.0.0.1      InLoopBack0
     10.0.12.0/24   OSPF    10   3124        D   10.0.14.1      Serial1/0/0
     10.0.14.0/24   Direct  0    0           D   10.0.14.4      Serial1/0/0
     10.0.14.1/32   Direct  0    0           D   10.0.14.1      Serial1/0/0
     10.0.14.4/32   Direct  0    0           D   127.0.0.1      InLoopBack0
   10.0.14.255/32   Direct  0    0           D   127.0.0.1      InLoopBack0
     10.0.15.0/24   OSPF    10   1563        D   10.0.14.1      Serial1/0/0
      127.0.0.0/8   Direct  0    0           D   127.0.0.1      InLoopBack0
     127.0.0.1/32   Direct  0    0           D   127.0.0.1      InLoopBack0
127.255.255.255/32  Direct  0    0           D   127.0.0.1      InLoopBack0
255.255.255.255/32  Direct  0    0           D   127.0.0.1      InLoopBack0

[R5]display ip routing-table
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Routing Tables: Public
         Destinations : 16      Routes : 16
```

```
Destination/Mask    Proto  Pre  Cost      Flags NextHop         Interface

      10.0.1.1/32   OSPF   10   1           D   10.0.15.1   GigabitEthernet0/0/0
      10.0.4.4/32   OSPF   10   1563        D   10.0.15.1   GigabitEthernet0/0/0
      10.0.5.5/32   Direct  0   0           D   127.0.0.1       InLoopBack0
     10.0.12.0/24   OSPF   10   1563        D   10.0.15.1   GigabitEthernet0/0/0
     10.0.14.0/24   OSPF   10   1563        D   10.0.15.1   GigabitEthernet0/0/0
     10.0.15.0/24   Direct  0   0           D   10.0.15.5   GigabitEthernet0/0/0
     10.0.15.5/32   Direct  0   0           D   127.0.0.1       InLoopBack0
   10.0.15.255/32   Direct  0   0           D   127.0.0.1       InLoopBack0
     10.0.35.0/24   Direct  0   0           D   10.0.35.5       Serial1/0/0
     10.0.35.3/32   Direct  0   0           D   10.0.35.3       Serial1/0/0
     10.0.35.5/32   Direct  0   0           D   127.0.0.1       InLoopBack0
   10.0.35.255/32   Direct  0   0           D   127.0.0.1       InLoopBack0
     127.0.0.0/8    Direct  0   0           D   127.0.0.1       InLoopBack0
     127.0.0.1/32   Direct  0   0           D   127.0.0.1       InLoopBack0
127.255.255.255/32  Direct  0   0           D   127.0.0.1       InLoopBack0
255.255.255.255/32  Direct  0   0           D   127.0.0.1       InLoopBack0
```

在R2，R3，R5上配置RIP。

R2的S2/0/0和Loopback0连接的网段运行RIP，

```
[R2]rip
[R2-rip-1]version 2
[R2-rip-1]network 10.0.0.0
```

R3的S2/0/0，S3/0/0和Loopback0连接的网段运行RIP，

```
[R3]rip
[R3-rip-1]version 2
[R3-rip-1]network 10.0.0.0
```

R5的S1/0/0和Loopback0连接的网段运行RIP，

```
[R5]rip
[R5-rip-1]version 2
[R5-rip-1]network 10.0.0.0
```

检查是否学习到其他设备上的Loopback0地址。

```
[R2]display ip routing-table
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Routing Tables: Public
       Destinations : 20      Routes : 20
```

```
Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface

      10.0.1.1/32   OSPF    10   1562        D   10.0.12.1       Serial1/0/0
      10.0.2.2/32   Direct  0    0           D   127.0.0.1       InLoopBack0
      10.0.3.3/32   RIP     100  1           D   10.0.23.3       Serial2/0/0
      10.0.4.4/32   OSPF    10   3124        D   10.0.12.1       Serial1/0/0
      10.0.5.5/32   RIP     100  2           D   10.0.23.3       Serial2/0/0
     10.0.12.0/24   Direct  0    0           D   10.0.12.2       Serial1/0/0
     10.0.12.1/32   Direct  0    0           D   10.0.12.1       Serial1/0/0
     10.0.12.2/32   Direct  0    0           D   127.0.0.1       InLoopBack0
   10.0.12.255/32   Direct  0    0           D   127.0.0.1       InLoopBack0
     10.0.14.0/24   OSPF    10   3124        D   10.0.12.1       Serial1/0/0
     10.0.15.0/24   OSPF    10   1563        D   10.0.12.1       Serial1/0/0
     10.0.23.0/24   Direct  0    0           D   10.0.23.2       Serial2/0/0
     10.0.23.2/32   Direct  0    0           D   127.0.0.1       InLoopBack0
     10.0.23.3/32   Direct  0    0           D   10.0.23.3       Serial2/0/0
   10.0.23.255/32   Direct  0    0           D   127.0.0.1       InLoopBack0
     10.0.35.0/24   RIP     100  1           D   10.0.23.3       Serial2/0/0
     127.0.0.0/8    Direct  0    0           D   127.0.0.1       InLoopBack0
     127.0.0.1/32   Direct  0    0           D   127.0.0.1       InLoopBack0
127.255.255.255/32  Direct  0    0           D   127.0.0.1       InLoopBack0
255.255.255.255/32  Direct  0    0           D   127.0.0.1       InLoopBack0

[R3]display ip routing-table
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Routing Tables: Public
         Destinations : 17       Routes : 17

Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface

      10.0.2.2/32   RIP     100  1           D   10.0.23.2       Serial2/0/0
      10.0.3.3/32   Direct  0    0           D   127.0.0.1       InLoopBack0
      10.0.5.5/32   RIP     100  1           D   10.0.35.5       Serial3/0/0
     10.0.12.0/24   RIP     100  1           D   10.0.23.2       Serial2/0/0
     10.0.15.0/24   RIP     100  1           D   10.0.35.5       Serial3/0/0
     10.0.23.0/24   Direct  0    0           D   10.0.23.3       Serial2/0/0
     10.0.23.2/32   Direct  0    0           D   10.0.23.2       Serial2/0/0
     10.0.23.3/32   Direct  0    0           D   127.0.0.1       InLoopBack0
   10.0.23.255/32   Direct  0    0           D   127.0.0.1       InLoopBack0
     10.0.35.0/24   Direct  0    0           D   10.0.35.3       Serial3/0/0
     10.0.35.3/32   Direct  0    0           D   127.0.0.1       InLoopBack0
```

```
      10.0.35.5/32   Direct  0    0           D   10.0.35.5       Serial3/0/0
    10.0.35.255/32   Direct  0    0           D   127.0.0.1       InLoopBack0
      127.0.0.0/8    Direct  0    0           D   127.0.0.1       InLoopBack0
      127.0.0.1/32   Direct  0    0           D   127.0.0.1       InLoopBack0
127.255.255.255/32   Direct  0    0           D   127.0.0.1       InLoopBack0
255.255.255.255/32   Direct  0    0           D   127.0.0.1       InLoopBack0

[R5]display ip routing-table
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Routing Tables: Public
        Destinations : 19      Routes : 19

Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface

       10.0.1.1/32   OSPF    10   1           D   10.0.15.1   GigabitEthernet0/0/0
       10.0.2.2/32   RIP     100  2           D   10.0.35.3       Serial1/0/0
       10.0.3.3/32   RIP     100  1           D   10.0.35.3       Serial1/0/0
       10.0.4.4/32   OSPF    10   1563        D   10.0.15.1   GigabitEthernet0/0/0
       10.0.5.5/32   Direct  0    0           D   127.0.0.1       InLoopBack0
      10.0.12.0/24   OSPF    10   1563        D   10.0.15.1   GigabitEthernet0/0/0
      10.0.14.0/24   OSPF    10   1563        D   10.0.15.1   GigabitEthernet0/0/0
      10.0.15.0/24   Direct  0    0           D   10.0.15.5   GigabitEthernet0/0/0
      10.0.15.5/32   Direct  0    0           D   127.0.0.1       InLoopBack0
    10.0.15.255/32   Direct  0    0           D   127.0.0.1       InLoopBack0
      10.0.23.0/24   RIP     100  1           D   10.0.35.3       Serial1/0/0
      10.0.35.0/24   Direct  0    0           D   10.0.35.5       Serial1/0/0
      10.0.35.3/32   Direct  0    0           D   10.0.35.3       Serial1/0/0
      10.0.35.5/32   Direct  0    0           D   127.0.0.1       InLoopBack0
    10.0.35.255/32   Direct  0    0           D   127.0.0.1       InLoopBack0
      127.0.0.0/8    Direct  0    0           D   127.0.0.1       InLoopBack0
      127.0.0.1/32   Direct  0    0           D   127.0.0.1       InLoopBack0
127.255.255.255/32   Direct  0    0           D   127.0.0.1       InLoopBack0
255.255.255.255/32   Direct  0    0           D   127.0.0.1       InLoopBack0
```

步骤三. 配置前缀列表过滤路由信息

在R1上创建静态路由1.1.1.1/32，1.1.1.0/24，1.1.1.0/25，1.1.0.0/16，1.0.0.0/8分别指向NULL 0接口。并使用**import-route static**命令将这些路由引入OSPF中。

```
[R1]ip route-static 1.1.1.1 255.255.255.255 NULL 0
```

```
[R1]ip route-static 1.1.1.0 255.255.255.0 NULL 0
[R1]ip route-static 1.1.1.0 255.255.255.128 NULL 0
[R1]ip route-static 1.1.0.0 255.255.0.0 NULL 0
[R1]ip route-static 1.0.0.0 255.0.0.0 NULL 0
[R1]ospf 1
[R1-ospf-1]import-route static
```

在R4上检查是否接收到R1上添加的这几条静态路由。

```
[R4]display ip routing-table
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Routing Tables: Public
        Destinations : 17       Routes : 17

Destination/Mask    Proto  Pre  Cost     Flags NextHop       Interface

       1.0.0.0/8     O_ASE   150  1           D   10.0.14.1       Serial1/0/0
       1.1.0.0/16    O_ASE   150  1           D   10.0.14.1       Serial1/0/0
       1.1.1.0/24    O_ASE   150  1           D   10.0.14.1       Serial1/0/0
       1.1.1.0/25    O_ASE   150  1           D   10.0.14.1       Serial1/0/0
       1.1.1.1/32    O_ASE   150  1           D   10.0.14.1       Serial1/0/0
      10.0.1.1/32    OSPF    10   1562        D   10.0.14.1       Serial1/0/0
      10.0.4.4/32    Direct  0    0           D   127.0.0.1       InLoopBack0
     10.0.12.0/24    OSPF    10   3124        D   10.0.14.1       Serial1/0/0
     10.0.14.0/24    Direct  0    0           D   10.0.14.4       Serial1/0/0
     10.0.14.1/32    Direct  0    0           D   10.0.14.1       Serial1/0/0
     10.0.14.4/32    Direct  0    0           D   127.0.0.1       InLoopBack0
   10.0.14.255/32    Direct  0    0           D   127.0.0.1       InLoopBack0
     10.0.15.0/24    OSPF    10   1563        D   10.0.14.1       Serial1/0/0
     127.0.0.0/8     Direct  0    0           D   127.0.0.1       InLoopBack0
     127.0.0.1/32    Direct  0    0           D   127.0.0.1       InLoopBack0
127.255.255.255/32   Direct  0    0           D   127.0.0.1       InLoopBack0
255.255.255.255/32   Direct  0    0           D   127.0.0.1       InLoopBack0
```

在R1上配置前缀列表pref_r1，匹配路由1.1.1.0/24。

```
[R1]ip ip-prefix pref_r1 index 10 permit 1.1.1.0 24 greater-equal 24 less-equal 24
```

创建路由策略policy_r1，调用前缀列表pref_r1，控制R1上引入的静态路由信息。

```
[R1]route-policy policy_r1 per node 10
[R1-route-policy]if-match ip-prefix pref_r1
```

```
[R1-route-policy]ospf
[R1-ospf-1]import-route static route-policy policy_r1
```

在R4上查看路由表。

```
[R4]display ip routing-table
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Routing Tables: Public
        Destinations : 13      Routes : 13

Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface

       1.1.1.0/24   O_ASE   150  1           D   10.0.14.1       Serial1/0/0
      10.0.1.1/32   OSPF    10   1562        D   10.0.14.1       Serial1/0/0
      10.0.4.4/32   Direct  0    0           D   127.0.0.1       InLoopBack0
     10.0.12.0/24   OSPF    10   3124        D   10.0.14.1       Serial1/0/0
     10.0.14.0/24   Direct  0    0           D   10.0.14.4       Serial1/0/0
     10.0.14.1/32   Direct  0    0           D   10.0.14.1       Serial1/0/0
     10.0.14.4/32   Direct  0    0           D   127.0.0.1       InLoopBack0
   10.0.14.255/32   Direct  0    0           D   127.0.0.1       InLoopBack0
     10.0.15.0/24   OSPF    10   1563        D   10.0.14.1       Serial1/0/0
     127.0.0.0/8    Direct  0    0           D   127.0.0.1       InLoopBack0
     127.0.0.1/32   Direct  0    0           D   127.0.0.1       InLoopBack0
127.255.255.255/32  Direct  0    0           D   127.0.0.1       InLoopBack0
255.255.255.255/32  Direct  0    0           D   127.0.0.1       InLoopBack0
```

步骤四. 使用 Route-policy 过滤路由信息避免路由环路

在R4上创建Loopback1，地址为10.1.4.4/24，使用**import-route direct**命令引入OSPF中。

```
[R4]interface LoopBack 1
[R4-LoopBack1]ip address 10.1.4.4 255.255.255.0
[R4-LoopBack1]ospf 1
[R4-ospf-1]import-route direct
```

在R2上将OSPF引入到RIP中，在R5上将RIP引入到OSPF中。

```
[R2]rip
[R2-rip-1]import-route ospf

[R5]ospf
```

```
[R5-ospf-1]import-route rip
```

在R1上测试到10.1.4.4的连通性。

```
[R1]ping 10.1.4.4
  PING 10.1.4.4: 56  data bytes, press CTRL_C to break
    Request time out
    Request time out
    Request time out
    Request time out
    Request time out

  --- 10.1.4.4 ping statistics ---
    5 packet(s) transmitted
    0 packet(s) received
    100.00% packet loss
```

发现不通，查看R1路由表。

```
[R1]display ip routing-table
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Routing Tables: Public
         Destinations : 28       Routes : 28

Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface

        1.0.0.0/8   Static  60   0           D   0.0.0.0         NULL0
       1.1.0.0/16   Static  60   0           D   0.0.0.0         NULL0
       1.1.1.0/24   Static  60   0           D   0.0.0.0         NULL0
       1.1.1.0/25   Static  60   0           D   0.0.0.0         NULL0
       1.1.1.1/32   Static  60   0           D   0.0.0.0         NULL0
      10.0.1.1/32   Direct  0    0           D   127.0.0.1       InLoopBack0
      10.0.2.2/32   O_ASE   150  1           D   10.0.15.5   GigabitEthernet0/0/0
      10.0.3.3/32   O_ASE   150  1           D   10.0.15.5   GigabitEthernet0/0/0
      10.0.4.4/32   OSPF    10   1562        D   10.0.14.4       Serial3/0/0
      10.0.5.5/32   O_ASE   150  1           D   10.0.15.5   GigabitEthernet0/0/0
     10.0.12.0/24   Direct  0    0           D   10.0.12.1       Serial1/0/0
     10.0.12.1/32   Direct  0    0           D   127.0.0.1       InLoopBack0
     10.0.12.2/32   Direct  0    0           D   10.0.12.2       Serial1/0/0
   10.0.12.255/32   Direct  0    0           D   127.0.0.1       InLoopBack0
     10.0.14.0/24   Direct  0    0           D   10.0.14.1       Serial3/0/0
     10.0.14.1/32   Direct  0    0           D   127.0.0.1       InLoopBack0
     10.0.14.4/32   Direct  0    0           D   10.0.14.4       Serial3/0/0
```

```
   10.0.14.255/32  Direct  0    0           D   127.0.0.1       InLoopBack0
     10.0.15.0/24  Direct  0    0           D   10.0.15.1   GigabitEthernet0/0/0
     10.0.15.1/32  Direct  0    0           D   127.0.0.1       InLoopBack0
   10.0.15.255/32  Direct  0    0           D   127.0.0.1       InLoopBack0
     10.0.23.0/24  O_ASE   150  1           D   10.0.15.5   GigabitEthernet0/0/0
     10.0.35.0/24  O_ASE   150  1           D   10.0.15.5   GigabitEthernet0/0/0
      10.1.4.0/24  O_ASE   150  1           D   10.0.15.5   GigabitEthernet0/0/0
      127.0.0.0/8  Direct  0    0           D   127.0.0.1       InLoopBack0
     127.0.0.1/32  Direct  0    0           D   127.0.0.1       InLoopBack0
127.255.255.255/32 Direct  0    0           D   127.0.0.1       InLoopBack0
255.255.255.255/32 Direct  0    0           D   127.0.0.1       InLoopBack0
```

此时R1上路由10.1.4.0/24的下一跳是R5。

再分别查看R2，R3，R5上路由表中的10.1.4.0/24这条路由。

```
[R2]display ip routing-table 10.1.4.0
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Routing Table : Public
Summary Count : 1
Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface

      10.1.4.0/24   O_ASE   150  1           D   10.0.12.1       Serial1/0/0

[R3]display ip routing-table 10.1.4.0
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Routing Table : Public
Summary Count : 1
Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface

      10.1.4.0/24   RIP     100  1           D   10.0.23.2       Serial2/0/0

[R5]display ip routing-table 10.1.4.0
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Routing Table : Public
Summary Count : 1
Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface

      10.1.4.0/24   RIP     100  2           D   10.0.35.3       Serial1/0/0
```

在R1上使用**tracert**命令查看到地址10.1.4.4的路径。

```
[R1]tracert 10.1.4.4
 traceroute to  10.1.4.4(10.1.4.4), max hops: 30 ,packet length: 40,press CTRL_C
to break
 1 10.0.15.5 61 ms  2 ms  2 ms
 2 10.0.35.3 29 ms  28 ms  29 ms
 3 10.0.23.2 31 ms  36 ms  36 ms
 4 10.0.12.1 34 ms  36 ms  36 ms
 5 10.0.15.5 34 ms  37 ms  37 ms
 6 10.0.35.3 55 ms  59 ms  59 ms
 7 10.0.23.2 60 ms  66 ms  66 ms
 8 10.0.12.1 63 ms  66 ms  66 ms
 9 10.0.15.5 65 ms  67 ms  67 ms
```

此时发现环路。

原因是配置了路由引入后，R5既可以在RIP域中学习到10.1.4.0/24，也可以在OSPF域中学习到这条路由。

由于RIP的路由优先级比OSPF外部路由的优先级高，从而导致R5使用了从RIP域中学习到的路由。

R1能同时从R5和R4学习到这条路由。对于这两条外部路由，同是OSPF外部路由，所以比较二者的cost值。由于R1连接R5的是千兆链路，优于R1连接R4的串行链路，使R1选择使用从R5学习的路由，最终导致环路。

在R5上使用路由策略policy_r5，给路由10.1.4.0/24打上标记100。

```
[R5]acl number 2001
[R5-acl-basic-2001]rule 0 permit source 10.1.4.0 0.0.0.255
[R5-acl-basic-2001]route-policy add_tag per node 10
[R5-route-policy]if-match acl 2001
[R5-route-policy]apply tag 100
[R5-route-policy]route-policy add_tag permit node 20
[R5-route-policy]ospf
[R5-ospf-1]import-route rip route-policy add_tag
```

在R1上查看OSPF路由信息。

```
[R1]display ospf routing

         OSPF Process 1 with Router ID 10.0.12.1
                Routing Tables

 Routing for Network
 Destination        Cost  Type       NextHop         AdvRouter       Area
 10.0.1.1/32        0     Stub       10.0.1.1        10.0.12.1       0.0.0.0
```

```
 10.0.12.0/24       1562 Stub       10.0.12.1       10.0.12.1       0.0.0.0
 10.0.14.0/24       1562 Stub       10.0.14.1       10.0.12.1       0.0.0.0
 10.0.15.0/24       1    Transit    10.0.15.1       10.0.12.1       0.0.0.0
 10.0.4.4/32        1562 Stub       10.0.14.4       10.0.14.4       0.0.0.0

 Routing for ASEs
 Destination        Cost      Type       Tag         NextHop        AdvRouter
 1.1.1.0/24         1         Type2      1           10.0.15.5      10.0.35.5
 10.0.2.2/32        1         Type2      1           10.0.15.5      10.0.35.5
 10.0.3.3/32        1         Type2      1           10.0.15.5      10.0.35.5
 10.0.5.5/32        1         Type2      1           10.0.15.5      10.0.35.5
 10.0.14.1/32       1         Type2      1           10.0.15.5      10.0.35.5
 10.0.23.0/24       1         Type2      1           10.0.15.5      10.0.35.5
 10.0.35.0/24       1         Type2      1           10.0.15.5      10.0.35.5
 10.1.4.0/24        1         Type2      100         10.0.15.5      10.0.35.5

 Total Nets: 13
 Intra Area: 5  Inter Area: 0  ASE: 8  NSSA: 0
```

在R1上发现携带标记100的10.1.4.0/24路由，再次R1验证是从R5获得这条路由。

为了解决环路问题，在R5上向OSPF引入RIP路由时，过滤掉路由10.1.4.0/24即可。

在R5上配置路由策略route_delete，控制向OSPF引入的RIP路由。

```
[R5]route-policy route_delete deny node 10
[R5-route-policy]if-match acl 2001
[R5-route-policy]route-policy route_delete permit node 20
[R5-route-policy]ospf 1
[R5-ospf-1]import-route rip route-policy route_delete
```

在R1上查看路由表，

```
[R1]display ip routing-table
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Routing Tables: Public
         Destinations : 28       Routes : 28

Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface

        1.0.0.0/8    Static  60   0           D   0.0.0.0         NULL0
        1.1.0.0/16   Static  60   0           D   0.0.0.0         NULL0
```

```
      1.1.1.0/24   Static  60   0          D   0.0.0.0         NULL0
      1.1.1.0/25   Static  60   0          D   0.0.0.0         NULL0
      1.1.1.1/32   Static  60   0          D   0.0.0.0         NULL0
     10.0.1.1/32   Direct  0    0          D   127.0.0.1       InLoopBack0
     10.0.2.2/32   O_ASE   150  1          D   10.0.15.5   GigabitEthernet0/0/0
     10.0.3.3/32   O_ASE   150  1          D   10.0.15.5   GigabitEthernet0/0/0
     10.0.4.4/32   OSPF    10   1562       D   10.0.14.4       Serial3/0/0
     10.0.5.5/32   O_ASE   150  1          D   10.0.15.5   GigabitEthernet0/0/0
    10.0.12.0/24   Direct  0    0          D   10.0.12.1       Serial1/0/0
    10.0.12.1/32   Direct  0    0          D   127.0.0.1       InLoopBack0
    10.0.12.2/32   Direct  0    0          D   10.0.12.2       Serial1/0/0
  10.0.12.255/32   Direct  0    0          D   127.0.0.1       InLoopBack0
    10.0.14.0/24   Direct  0    0          D   10.0.14.1       Serial3/0/0
    10.0.14.1/32   Direct  0    0          D   127.0.0.1       InLoopBack0
    10.0.14.4/32   Direct  0    0          D   10.0.14.4       Serial3/0/0
  10.0.14.255/32   Direct  0    0          D   127.0.0.1       InLoopBack0
    10.0.15.0/24   Direct  0    0          D   10.0.15.1   GigabitEthernet0/0/0
    10.0.15.1/32   Direct  0    0          D   127.0.0.1       InLoopBack0
  10.0.15.255/32   Direct  0    0          D   127.0.0.1       InLoopBack0
    10.0.23.0/24   O_ASE   150  1          D   10.0.15.5   GigabitEthernet0/0/0
    10.0.35.0/24   O_ASE   150  1          D   10.0.15.5   GigabitEthernet0/0/0
     10.1.4.0/24   O_ASE   150  1          D   10.0.14.4       Serial3/0/0
    127.0.0.0/8    Direct  0    0          D   127.0.0.1       InLoopBack0
    127.0.0.1/32   Direct  0    0          D   127.0.0.1       InLoopBack0
127.255.255.255/32 Direct  0    0          D   127.0.0.1       InLoopBack0
255.255.255.255/32 Direct  0    0          D   127.0.0.1       InLoopBack0
```

显示路由信息正常。

测试R1与地址10.1.4.4之间的连通性。

```
[R1]ping 10.1.4.4
  PING 10.1.4.4: 56  data bytes, press CTRL_C to break
    Reply from 10.1.4.4: bytes=56 Sequence=1 ttl=255 time=33 ms
    Reply from 10.1.4.4: bytes=56 Sequence=2 ttl=255 time=29 ms
    Reply from 10.1.4.4: bytes=56 Sequence=3 ttl=255 time=29 ms
    Reply from 10.1.4.4: bytes=56 Sequence=4 ttl=255 time=29 ms
    Reply from 10.1.4.4: bytes=56 Sequence=5 ttl=255 time=29 ms

  --- 10.1.4.4 ping statistics ---
    5 packet(s) transmitted
    5 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 29/29/33 ms
```

在R1上检测到达地址10.1.4.4的路径。

```
[R1]tracert 10.1.4.4
 traceroute to  10.1.4.4(10.1.4.4), max hops: 30 ,packet length: 40,press CTRL_C
to break
 1 10.0.14.4 61 ms  29 ms  29 ms
```

此时环路问题解决，R1上路由10.1.4.0/24的下一跳为R4。

步骤五.　使用 Route-policy 修改路由优先级避免路由环路

查看R5的路由表。观察路由10.1.4.0/24的下一跳。

```
[R5]display ip routing-table
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Routing Tables: Public
        Destinations : 22      Routes : 22

Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface

       1.1.1.0/24   RIP     100  2           D   10.0.35.3       Serial1/0/0
      10.0.1.1/32   OSPF    10   1           D   10.0.15.1   GigabitEthernet0/0/0
      10.0.2.2/32   RIP     100  2           D   10.0.35.3       Serial1/0/0
      10.0.3.3/32   RIP     100  1           D   10.0.35.3       Serial1/0/0
      10.0.4.4/32   OSPF    10   1563         D   10.0.15.1  GigabitEthernet0/0/0
      10.0.5.5/32   Direct  0    0           D   127.0.0.1       InLoopBack0
     10.0.12.0/24   OSPF    10   1563         D   10.0.15.1  GigabitEthernet0/0/0
     10.0.14.0/24   OSPF    10   1563         D   10.0.15.1  GigabitEthernet0/0/0
     10.0.14.1/32   RIP     100  2           D   10.0.35.3       Serial1/0/0
     10.0.15.0/24   Direct  0    0           D   10.0.15.5   GigabitEthernet0/0/0
     10.0.15.5/32   Direct  0    0           D   127.0.0.1       InLoopBack0
   10.0.15.255/32   Direct  0    0           D   127.0.0.1       InLoopBack0
     10.0.23.0/24   RIP     100  1           D   10.0.35.3       Serial1/0/0
     10.0.35.0/24   Direct  0    0           D   10.0.35.5       Serial1/0/0
     10.0.35.3/32   Direct  0    0           D   10.0.35.3       Serial1/0/0
     10.0.35.5/32   Direct  0    0           D   127.0.0.1       InLoopBack0
   10.0.35.255/32   Direct  0    0           D   127.0.0.1       InLoopBack0
      10.1.4.0/24   RIP     100  2           D   10.0.35.3       Serial1/0/0
     127.0.0.0/8    Direct  0    0           D   127.0.0.1       InLoopBack0
     127.0.0.1/32   Direct  0    0           D   127.0.0.1       InLoopBack0
127.255.255.255/32  Direct  0    0           D   127.0.0.1       InLoopBack0
```

```
255.255.255.255/32  Direct  0    0          D   127.0.0.1      InLoopBack0
```

在R5上检测到达地址10.1.4.4的路径。

```
[R5]tracert 10.1.4.4
 traceroute to  10.1.4.4(10.1.4.4), max hops: 30 ,packet length: 40,press CTRL_C
to break
 1 10.0.35.3 62 ms  24 ms  24 ms
 2 10.0.23.2 43 ms  44 ms  44 ms
 3 10.0.12.1 33 ms  33 ms  33 ms
 4 10.0.14.4 74 ms  55 ms  55 ms
```

发现通过路由过滤的方法虽然能解决环路问题，但是R5仍旧是从RIP域学习到路由10.1.4.0/24。显然次优路径问题仍然没有得到解决。

为了解决环路问题以及次优路径问题，阻止R5从RIP域中获得路由10.1.4.0/24，使R5从OSPF域学习到路由10.1.4.0/24。

删除R5上的策略route_delete。

```
[R5]undo route-policy route_delete
```

在R5上配置策略route_pref，将10.1.4.0/24的路由优先级修改为180，使其小于OSPF外部路由的路由优先级。

```
[R5]route-policy route_pref per node 10
[R5-route-policy]if-match acl 2001
[R5-route-policy]apply preference 180
```

使用路由策略route_pref来控制RIP发布给OSPF的路由。

```
[R5]rip
[R5-rip-1]preference route-policy route_pref
```

检查R5，R1的路由表，观察路由10.1.4.0/24的下一跳。

```
[R5]display ip routing-table 10.1.4.0
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Routing Table : Public
Summary Count : 1
Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface

      10.1.4.0/24   O_ASE   150  1           D   10.0.15.1   GigabitEthernet0/0/0
```

```
[R1]display ip routing-table 10.1.4.0
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Routing Table : Public
Summary Count : 1
Destination/Mask    Proto   Pre  Cost      Flags NextHop         Interface

      10.1.4.0/24   O_ASE   150  1           D   10.0.14.4       Serial3/0/0
```

检查R1与地址10.1.4.4的连通性。

```
[R1]ping 10.1.4.4
  PING 10.1.4.4: 56  data bytes, press CTRL_C to break
    Reply from 10.1.4.4: bytes=56 Sequence=1 ttl=255 time=39 ms
    Reply from 10.1.4.4: bytes=56 Sequence=2 ttl=255 time=35 ms
    Reply from 10.1.4.4: bytes=56 Sequence=3 ttl=255 time=35 ms
    Reply from 10.1.4.4: bytes=56 Sequence=4 ttl=255 time=35 ms
    Reply from 10.1.4.4: bytes=56 Sequence=5 ttl=255 time=35 ms

  --- 10.1.4.4 ping statistics ---
    5 packet(s) transmitted
    5 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 35/35/39 ms
```

在R1上检测到达地址10.1.4.4的路径。

```
[R1]tracert 10.1.4.4
 traceroute to  10.1.4.4(10.1.4.4), max hops: 30 ,packet length: 40,press CTRL_C
to break
 1 10.0.14.4 61 ms  25 ms  25 ms
```

在R5上检测到达地址10.1.4.4的路径。

```
[R5]tracert 10.1.4.4
 traceroute to  10.1.4.4(10.1.4.4), max hops: 30 ,packet length: 40,press CTRL_C
to break
 1 10.0.15.1 61 ms  2 ms  2 ms
 2 10.0.14.4 41 ms  28 ms  27 ms
```

此时环路问题解决。

在R1上，路由10.1.4.0/24的下一跳是R4。在R5上，路由10.1.4.0/24的下一跳是R1。次优路由也得到解决。

附加实验：思考并验证

在步骤三中，思考通过访问控制列表是否可以实现同样的效果，它和前缀列表的区别又是什么。

在步骤五中，思考为什么R3路由表中的10.0.15.0/24出现有两个下一跳，而10.0.12.0/24只有一个下一跳。

最终设备配置

```
<R1>display current-configuration
[V200R001C00SPC200]
#
 sysname R1
#
interface Serial1/0/0
 link-protocol ppp
 ip address 10.0.12.1 255.255.255.0
#
interface Serial3/0/0
 link-protocol ppp
 ip address 10.0.14.1 255.255.255.0
#
interface GigabitEthernet0/0/0
 ip address 10.0.15.1 255.255.255.0
#
interface LoopBack0
 ip address 10.0.1.1 255.255.255.255
#
ospf 1
 import-route static route-policy policy_r1
 area 0.0.0.0
  network 10.0.12.1 0.0.0.0
  network 10.0.15.1 0.0.0.0
  network 10.0.14.1 0.0.0.0
  network 10.0.1.1 0.0.0.0
#
route-policy policy_r1 permit node 10
 if-match ip-prefix pref_r1
```

```
#
 ip ip-prefix pref_r1 index 10 permit 1.1.1.0 24 greater-equal 24 less-equal 24
#
 ip route-static 1.0.0.0 255.0.0.0 NULL0
 ip route-static 1.1.0.0 255.255.0.0 NULL0
 ip route-static 1.1.1.0 255.255.255.0 NULL0
 ip route-static 1.1.1.0 255.255.255.128 NULL0
 ip route-static 1.1.1.1 255.255.255.255 NULL0
#
return

<R2>display current-configuration
[V200R001C00SPC200]
#
 sysname R2
#
interface Serial1/0/0
 link-protocol ppp
 ip address 10.0.12.2 255.255.255.0
#
interface Serial2/0/0
 link-protocol ppp
 ip address 10.0.23.2 255.255.255.0
#
interface LoopBack0
 ip address 10.0.2.2 255.255.255.255
#
ospf 1
 area 0.0.0.0
  network 10.0.12.2 0.0.0.0
#
rip 1
 version 2
 network 10.0.0.0
 import-route ospf 1
#
return

<R3>display current-configuration
[V200R001C00SPC200]
#
 sysname R3
#
```

```
interface Serial2/0/0
 link-protocol ppp
 ip address 10.0.23.3 255.255.255.0
#
interface Serial3/0/0
 link-protocol ppp
 ip address 10.0.35.3 255.255.255.0
#
interface LoopBack0
 ip address 10.0.3.3 255.255.255.255
#
rip 1
 version 2
 network 10.0.0.0
#
return

<R4>display current-configuration
[V200R001C00SPC200]
#
 sysname R4
#
interface Serial1/0/0
 link-protocol ppp
 ip address 10.0.14.4 255.255.255.0
#
interface LoopBack0
 ip address 10.0.4.4 255.255.255.255
#
interface LoopBack1
 ip address 10.1.4.4 255.255.255.0
#
ospf 1
 import-route direct
 area 0.0.0.0
  network 10.0.14.4 0.0.0.0
  network 10.0.4.4 0.0.0.0
#
return

<R5>display current-configuration
[V200R001C00SPC200]
#
```

```
 sysname R5
#
interface Serial1/0/0
 link-protocol ppp
 ip address 10.0.35.5 255.255.255.0
#
interface GigabitEthernet0/0/0
 ip address 10.0.15.5 255.255.255.0
#
interface LoopBack0
 ip address 10.0.5.5 255.255.255.255
#
ospf 1
 import-route rip 1 route-policy route_delete
 area 0.0.0.0
  network 10.0.15.5 0.0.0.0
#
rip 1
 version 2
 network 10.0.0.0
 preference route-policy route_pref
#
route-policy add_tag permit node 10
 if-match acl 2001
 apply tag 100
#
route-policy add_tag permit node 20
#
route-policy route_pref permit node 10
 if-match acl 2001
 apply preference 180
#
return
```

第五章 组播协议

实验 5-1 组播、IGMP 及 PIM DM 协议

学习目的

- 掌握路由器启动组播路由功能的配置方法
- 掌握配置接口IGMP功能的方法
- 掌握PIM DM的配置方法
- 掌握查看和测试组播的方法
- 掌握PIM一些高级特性的配置方法

拓扑图

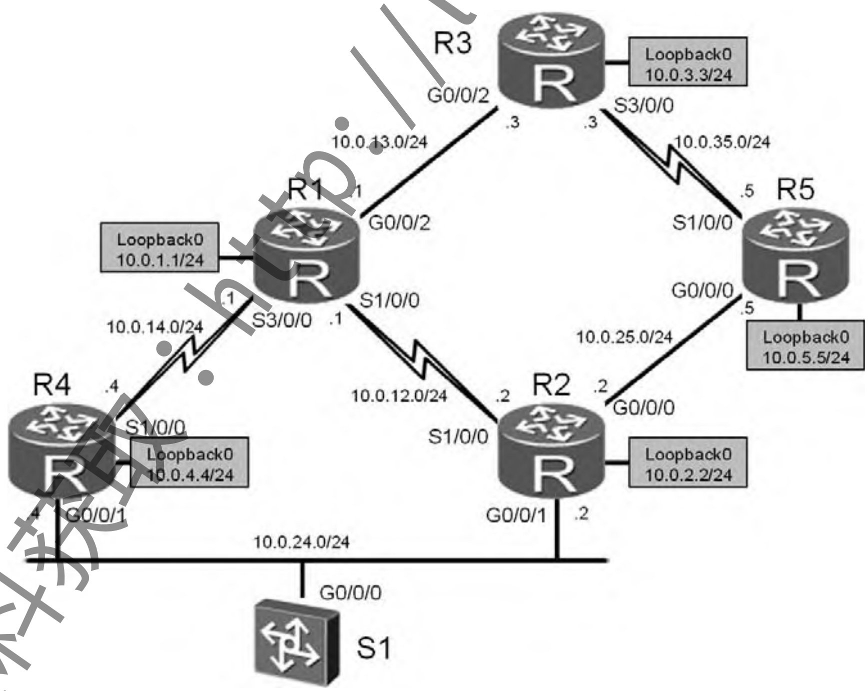

图5-1 组播、IGMP及PIM DM协议

场景

你是公司的网络管理员。公司准备使用组播来进行一些业务的转发。在当前网络上，网络规模较小，你决定使用PIM的DM模式来实现组播路由信息的学习。组播转发的实现过程中，你需要考虑到主机应用对IGMP不同版本的兼容，同时需要考虑使用合适的方式测试网络中组播是否正常工作。为了提升网络的效率和安全性，你采用了PIM DM的一些手段，包括PIM邻居的控制、嫁接和其他安全措施。同时实现网络的组播转发之前，你也遇到了一些网络故障，经过一些故障排除步骤，最终网络正常工作了。

学习任务

步骤一. 基础配置与 IP 编址

S2参与到本次实验（实现R1和R3的互联），但无需配置。实验之前，请清空S2的配置，并重启它。

给所有路由器配置IP地址和掩码。配置时注意所有的Loopback接口配置掩码均为32位。

```
<Huawei>system-view
Enter system view, return user view with Ctrl+Z.
[Huawei]sysname R1
[R1]interface Serial 1/0/0
[R1-Serial1/0/0]ip address 10.0.12.1 24
[R1-Serial1/0/0]interface Serial 3/0/0
[R1-Serial3/0/0]ip address 10.0.14.1 24
[R1-Serial3/0/0]interface Serial 1/0/0
[R1-Serial1/0/0]ip address 10.0.12.1 24
[R1-Serial1/0/0]interface GigabitEthernet 0/0/2
[R1-GigabitEthernet0/0/2]ip address 10.0.13.1 24
[R1-GigabitEthernet0/0/2]interface LoopBack 0
[R1-LoopBack0]ip address 10.0.1.1 24

<Huawei>system-view
Enter system view, return user view with Ctrl+Z.
[Huawei]sysname R2
[R2]interface Serial 1/0/0
[R2-Serial1/0/0]ip address 10.0.12.2 24
[R2-Serial1/0/0] interface GigabitEthernet 0/0/0
[R2-GigabitEthernet0/0/0]ip address 10.0.25.2 24
[R2-GigabitEthernet0/0/0]interface GigabitEthernet 0/0/1
```

```
[R2-GigabitEthernet0/0/1]ip address 10.0.24.2 24
[R2-GigabitEthernet0/0/1]interface LoopBack 0
[R2-LoopBack0]ip address 10.0.2.2 24

<Huawei>system-view
Enter system view, return user view with Ctrl+Z.
[Huawei]sysname R3
[R3]interface GigabitEthernet 0/0/2
[R3-GigabitEthernet0/0/2]ip address 10.0.13.3 24
[R3-GigabitEthernet0/0/2]interface Serial 3/0/0
[R3-Serial3/0/0]ip address 10.0.35.3 24
[R3-Serial3/0/0]interface LoopBack 0
[R3-LoopBack0]ip address 10.0.3.3 24

<Huawei>system-view
Enter system view, return user view with Ctrl+Z.
[Huawei]sysname R4
[R4]interface Serial 1/0/0
[R4-Serial1/0/0]ip address 10.0.14.4 24
[R4-Serial1/0/0]interface GigabitEthernet 0/0/1
[R4-GigabitEthernet0/0/1]ip address 10.0.24.4 24
[R4-GigabitEthernet0/0/1]interface LoopBack 0
[R4-LoopBack0]ip address 10.0.4.4 24

<Huawei>system-view
Enter system view, return user view with Ctrl+Z.
[Huawei]sysname R5
[R5]interface Serial 1/0/0
[R5-Serial1/0/0]ip address 10.0.35.5 24
[R5-Serial1/0/0]interface GigabitEthernet 0/0/0
[R5-GigabitEthernet0/0/0]ip address 10.0.25.5 24
[R5-GigabitEthernet0/0/0]interface LoopBack 0
[R5-LoopBack0]ip address 10.0.5.5 24
```

配置完成以后，验证路由器之间的连通性。

```
[R1]ping -c 1 10.0.13.3
  PING 10.0.13.3: 56  data bytes, press CTRL_C to break
    Reply from 10.0.13.3: bytes=56 Sequence=1 ttl=255 time=5 ms

  --- 10.0.13.3 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
```

```
    round-trip min/avg/max = 5/5/5 ms

[R1]ping -c 1 10.0.12.2
  PING 10.0.12.2: 56  data bytes, press CTRL_C to break
    Reply from 10.0.12.2: bytes=56 Sequence=1 ttl=255 time=37 ms

  --- 10.0.12.2 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 37/37/37 ms

[R1]ping -c 1 10.0.14.4
  PING 10.0.14.4: 56  data bytes, press CTRL_C to break
    Reply from 10.0.14.4: bytes=56 Sequence=1 ttl=255 time=38 ms

  --- 10.0.14.4 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 38/38/38 ms

[R5]ping -c 1 10.0.35.3
  PING 10.0.35.3: 56 data bytes, press CTRL_C to break
    Reply from 10.0.35.3: bytes=56 Sequence=1 ttl=255 time=33 ms

  --- 10.0.35.3 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 33/33/33 ms

[R5]ping -c 1 10.0.25.2
  PING 10.0.25.2: 56 data bytes, press CTRL_C to break
    Reply from 10.0.25.2: bytes=56 Sequence=1 ttl=255 time=10 ms

  --- 10.0.25.2 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 10/10/10 ms
```

步骤二. 配置所有路由器启用组播路由功能

启用R1、R2、R3、R4及R5的组播路由功能。要开启组播功能，首先在系统视图下运行命令**multicast routing-enable**。

默认情况下，VRP平台的组播功能是关闭的。无论要使用PIM还是IGMP都需要先在全局下开启组播功能。

```
[R1]multicast routing-enable
```

对于要运行PIM DM的接口，在接口视图下运行**pim dm**开启组播路由协议。

```
[R1]interface GigabitEthernet 0/0/2
[R1-GigabitEthernet0/0/2]pim dm
[R1-GigabitEthernet0/0/2]interface Serial 1/0/0
[R1-Serial1/0/0]pim dm
[R1-Serial1/0/0]interface Serial 3/0/0
[R1-Serial3/0/0]pim dm
```

在R2、R3、R4、R5上进行相同的配置，注意在路由器与路由器的互联接口上开启PIM DM的功能。

```
[R2]multicast routing-enable
[R2]interface Serial 1/0/0
[R2-Serial1/0/0]pim dm
[R2-Serial1/0/0]interface GigabitEthernet 0/0/0
[R2-GigabitEthernet0/0/0]pim dm

[R3]multicast routing-enable
[R3]interface GigabitEthernet 0/0/2
[R3-GigabitEthernet0/0/2]pim dm
[R3-GigabitEthernet0/0/2]interface Serial 3/0/0
[R3-Serial3/0/0]pim dm

[R4]multicast routing-enable
[R4]interface Serial 1/0/0
[R4-Serial1/0/0]pim dm
[R4-Serial1/0/0]interface GigabitEthernet 0/0/1
[R4-GigabitEthernet0/0/1]pim dm

[R5]multicast routing-enable
[R5]interface Serial 1/0/0
[R5-Serial1/0/0]pim dm
[R5-Serial1/0/0]interface  GigabitEthernet 0/0/0
```

```
[R5-GigabitEthernet0/0/0]pim dm
```

配置完成后，查看路由器PIM在接口上的运行状态。

```
[R1]display pim interface
 VPN-Instance: public net
 Interface          State NbrCnt HelloInt DR-Pri     DR-Address
 GE0/0/2              up    1      30        1        10.0.13.3
 S1/0/0               up    1      30        1        10.0.12.2
 S3/0/0               up    1      30        1        10.0.14.4
```

可以看到R1有3个接口运行了PIM，并且在每个接口上各有一个邻居（NbrCnt），同时我们还可以发现，在一个网段上接口IP地址较大的路由器将成为这个网段的DR。

查看R1上接口G0/0/2的PIM详细信息。

```
[R1]display pim interface GigabitEthernet 0/0/2 verbose
 VPN-Instance: public net
 Interface: GigabitEthernet0/0/2, 10.0.13.1
     PIM version: 2
     PIM mode: Dense
     PIM state: up
     PIM DR: 10.0.13.3
     PIM DR Priority (configured): 1
     PIM neighbor count: 1
     PIM hello interval: 30 s
     PIM LAN delay (negotiated): 500 ms
     PIM LAN delay (configured): 500 ms
     PIM hello override interval (negotiated): 2500 ms
     PIM hello override interval (configured): 2500 ms
     PIM Silent: disabled
     PIM neighbor tracking (negotiated): disabled
     PIM neighbor tracking (configured): disabled
     PIM generation ID: 0X5325911
     PIM require-GenID: disabled
     PIM hello hold interval: 105 s
     PIM assert hold interval: 180 s
     PIM triggered hello delay: 5 s
     PIM J/P interval: 60 s
     PIM J/P hold interval: 210 s
     PIM state-refresh processing: enabled
     PIM state-refresh interval: 60 s
     PIM graft retry interval: 3 s
     PIM state-refresh capability on link: capable
```

```
    PIM dr-switch-delay timer : not configured
    Number of routers on link not using DR priority: 0
    Number of routers on link not using LAN delay: 0
    Number of routers on link not using neighbor tracking: 2
    ACL of PIM neighbor policy: -
    ACL of PIM ASM join policy: -
    ACL of PIM SSM join policy: -
    ACL of PIM join policy: -
```

可以看到PIM DM默认的Hello间隔是30秒，Hello的保持时间是Hello间隔的3.5倍，也就是105秒。

查看R1的邻居列表，共有3台路由器与R1形成PIM邻居关系，邻居默认的DR优先级均为1。

```
[R1]display pim neighbor
 VPN-Instance: public net
 Total Number of Neighbors = 3

 Neighbor        Interface          Uptime   Expires  Dr-Priority
 10.0.13.3       GE0/0/2             01:40:27 00:01:18    1
 10.0.12.2       S1/0/0              01:42:21 00:01:24    1
 10.0.14.4       S3/0/0              01:38:02 00:01:16    1
```

查看邻居R3的详细信息，Uptime表示邻居关系已经建立的时间，Expiry time表示PIM邻居还有多少时间就要超时，LAN delay表示传递Prune剪枝消息的延迟时间，Override interval表示否决Prune剪枝的时间间隔。

```
[R1]display pim neighbor 10.0.13.3 verbose
 VPN-Instance: public net
 Neighbor: 10.0.13.3
    Interface: GigabitEthernet0/0/2
    Uptime: 01:41:00
    Expiry time: 00:01:45
    DR Priority: 1
    Generation ID: 0XD1A5CA9
    Holdtime: 105 s
    LAN delay: 500 ms
    Override interval: 2500 ms
    State refresh interval: 60 s
    Neighbor tracking: Disabled
```

步骤三. 配置 IGMP

在这个实验中，我们模拟组播用户连接在交换机S1上。在R2和R4的G0/0/1接口开启IGMP功能。要开启IGMP的功能，在接口模式下运行**igmp enable**。

```
[R2]interface GigabitEthernet 0/0/1
[R2-GigabitEthernet0/0/1]igmp enable

[R4]interface GigabitEthernet 0/0/1
[R4-GigabitEthernet0/0/1]igmp enable
```

为了实验需要，在R2和R4的G0/0/1接口添加静态组播组。这样，该接口始终会转发目的地址为225.1.1.1的组播流量。

```
[R2]interface GigabitEthernet 0/0/1
[R2-GigabitEthernet0/0/1]igmp static-group 225.1.1.1

[R4]interface GigabitEthernet 0/0/1
[R4-GigabitEthernet0/0/1]igmp static-group 225.1.1.1
```

默认情况下，VRP平台使用的IGMP版本为v2。从下面的输出中可以看到现在在G0/0/1接口所在网段的查询器是10.0.24.2，即为R2。对于IGMP v2来说，选取网段上IP地址较小的那台路由器作为查询器。

```
[R2]display igmp interface GigabitEthernet 0/0/1
Interface information
 GigabitEthernet0/0/1(10.0.24.2):
   IGMP is enabled
   Current IGMP version is 2
   IGMP state: up
   IGMP group policy: none
   IGMP limit: -
   Value of query interval for IGMP (negotiated): -
   Value of query interval for IGMP (configured): 60 s
   Value of other querier timeout for IGMP: 0 s
   Value of maximum query response time for IGMP: 10 s
   Querier for IGMP: 10.0.24.2 (this router)

[R4]display igmp interface GigabitEthernet 0/0/1
Interface information
 GigabitEthernet0/0/1(10.0.24.4):
   IGMP is enabled
   Current IGMP version is 2
   IGMP state: up
```

```
  IGMP group policy: none
  IGMP limit: -
  Value of query interval for IGMP (negotiated): -
  Value of query interval for IGMP (configured): 60 s
  Value of other querier timeout for IGMP: 123 s
  Value of maximum query response time for IGMP: 10 s
  Querier for IGMP: 10.0.24.2
```

查看接口的静态IGMP组，可以看到225.1.1.1这个组是刚才我们手工添加的组播组。

```
[R2]display igmp group static
Static join group information
 Total 1 entry, Total 1 active entry
  Group Address   Source Address  Interface          State    Expires
  225.1.1.1         0.0.0.0        GE0/0/1            UP       never
```

在接口下查看IGMP路由表。

```
[R2]display igmp routing-table
Routing table
 Total 1 entry

 00001. (*, 225.1.1.1)
      List of 1 downstream interface
       GigabitEthernet0/0/1 (10.0.24.2),
                 Protocol: STATIC
```

若在接口上只配置了IGMP，没有配置PIM，且接口为查询器的情况下，才会生成IGMP路由表项。该路由表条目在R4上是看不到的，因为现在R2是网段10.0.24.0/24的查询器。

默认情况下，查询器的查询周期为60秒，为了加快用户加入组播组的速度，我们可以通过**igmp timer query**修改发送查询报文的时间间隔。

```
[R2]interface GigabitEthernet 0/0/1
[R2-GigabitEthernet0/0/1]igmp timer query 20
```

配置完成以后，验证配置已生效。

```
[R2]display igmp interface GigabitEthernet 0/0/1
Interface information
 GigabitEthernet0/0/1(10.0.24.4):
  IGMP is enabled
  Current IGMP version is 1
  IGMP state: up
```

```
  IGMP group policy: none
  IGMP limit: -
  Value of query interval for IGMP (negotiated): -
  Value of query interval for IGMP (configured): 20 s
  Value of other querier timeout for IGMP: 0 s
  Value of maximum query response time for IGMP: -
  Querier for IGMP: 10.0.24.4 (this router)
```

开启Debugging后可以看到接口每隔20秒发送一次成员关系查询（general query）。

```
<R2>terminal debugging
<R2>debugging igmp query send
Dec 29 2011 16:33:17.350.1+00:00 R4 MGMD/7/QUERY:Send version 1 general query
on GigabitEthernet0/0/1(10.0.24.4) to destination 224.0.0.1 (G073088)
Dec 29 2011 16:33:37.130.1+00:00 R4 MGMD/7/QUERY:Send version 1 general query
on GigabitEthernet0/0/1(10.0.24.4) to destination 224.0.0.1 (G073088)
Dec 29 2011 16:33:57.510.1+00:00 R4 MGMD/7/QUERY:Send version 1 general query
on GigabitEthernet0/0/1(10.0.24.4) to destination 224.0.0.1 (G073088)
Dec 29 2011 16:34:17.480.1+00:00 R4 MGMD/7/QUERY:Send version 1 general query
on GigabitEthernet0/0/1(10.0.24.4) to destination 224.0.0.1 (G073088)
```

路由器的健壮系数描述了IGMP路由器的健壮程度。路由器默认的健壮系数为2，这里通过关闭接口的方式测试健壮系数。首先观察默认情况下IGMP查询消息的间隔。

```
<R2>terminal debugging
<R2>debugging igmp query send
Dec 31 2011 12:37:58.100.1+00:00 R2 MGMD/7/QUERY:Send version 2 general query
on GigabitEthernet0/0/1(10.0.24.2) to destination 224.0.0.1 (G073088)
Dec 31 2011 12:38:18.100.1+00:00 R2 MGMD/7/QUERY:Send version 2 general query
on GigabitEthernet0/0/1(10.0.24.2) to destination 224.0.0.1 (G073088)
<R2>system-view
Enter system view, return user view with Ctrl+Z.
Dec 31 2011 12:38:38.100.1+00:00 R2 MGMD/7/QUERY:Send version 2 general query
on GigabitEthernet0/0/1(10.0.24.2) to destination 224.0.0.1 (G073088)
[R2]interface GigabitEthernet 0/0/1
[R2-GigabitEthernet0/0/1]shutdown
Dec 31 2011 12:38:53+00:00 R2 %%01IFPDT/4/IF_STATE(l)[0]:Interface
GigabitEthernet0/0/1 has turned into DOWN state.
Dec 31 2011 12:38:53+00:00 R2 %%01IFNET/4/LINK_STATE(l)[1]:The line protocol  on
the interface GigabitEthernet0/0/1 has entered the DOWN state.
[R2-GigabitEthernet0/0/1]undo shutdown
Dec 31 2011 12:39:02+00:00 R2 %%01IFPDT/4/IF_STATE(l)[2]:Interface
```

```
GigabitEthernet0/0/1 has turned into UP state.
Dec 31 2011 12:39:02+00:00 R2 %%01IFNET/4/LINK_STATE(l)[3]:The line protocol  on
the interface GigabitEthernet0/0/1 has entered the UP state.
Dec 31 2011 12:39:03.100.1+00:00 R2 MGMD/7/QUERY:Send version 2 general query
on GigabitEthernet0/0/1(10.0.24.2) to destination 224.0.0.1 (G073088)
Dec 31 2011 12:39:08.100.1+00:00 R2 MGMD/7/QUERY:Send version 2 general query
on GigabitEthernet0/0/1(10.0.24.2) to destination 224.0.0.1 (G073088)
Dec 31 2011 12:39:28.100.1+00:00 R2 MGMD/7/QUERY:Send version 2 general query
on GigabitEthernet0/0/1(10.0.24.2) to destination 224.0.0.1 (G073088)
```

在没有关闭接口之前，路由器的接口仍按照每20秒一次的间隔发送普遍查询消息，当接口被关闭又重新打开之后，前面2个查询报文的时间间隔为5秒。当路由器启动时会发送“健壮系数”次的“普遍组查询消息”，发送间隔是“IGMP普遍组查询消息的发送间隔”的1/4。

执行命令**robust-count**可配置IGMP健壮系数，注意该参数只有在IGMP v2和IGMP v3中才有效，在R2的G0/0/1上将健壮系数修改为3。

```
[R2-GigabitEthernet0/0/1]igmp robust-count 3
```

再使用Debugging观察普遍查询消息的发送。

```
<R2>terminal debugging
Info: Current terminal debugging is on.
<R2>debugging igmp query send
Dec 31 2011 13:17:48.440.1+00:00 R2 MGMD/7/QUERY:Send version 2 general query
on GigabitEthernet0/0/1(10.0.24.2) to destination 224.0.0.1 (G073088)
<R2>system-view
Enter system view, return user view with Ctrl+Z.
[R2]interface GigabitEthernet 0/0/1
[R2-GigabitEthernet0/0/1]shutdown
Dec 31 2011 13:17:58+00:00 R2 %%01IFPDT/4/IF_STATE(l)[0]:Interface
GigabitEthernet0/0/1 has turned into DOWN state.
Dec 31 2011 13:17:58+00:00 R2 %%01IFNET/4/LINK_STATE(l)[1]:The line protocol  on
the interface GigabitEthernet0/0/1 has entered the DOWN state.
[R2-GigabitEthernet0/0/1]undo shutdown
Dec 31 2011 13:18:05+00:00 R2 %%01IFPDT/4/IF_STATE(l)[2]:Interface
GigabitEthernet0/0/1 has turned into UP state.
Dec 31 2011 13:18:05+00:00 R2 %%01IFNET/4/LINK_STATE(l)[3]:The line protocol  on
the interface GigabitEthernet0/0/1 has entered the UP state.
Dec 31 2011 13:18:06.440.1+00:00 R2 MGMD/7/QUERY:Send version 2 general query
on GigabitEthernet0/0/1(10.0.24.2) to destination 224.0.0.1 (G073088)
Dec 31 2011 13:18:11.440.1+00:00 R2 MGMD/7/QUERY:Send version 2 general query
on GigabitEthernet0/0/1(10.0.24.2) to destination 224.0.0.1 (G073088)
Dec 31 2011 13:18:16.440.1+00:00 R2 MGMD/7/QUERY:Send version 2 general query
```

```
on GigabitEthernet0/0/1(10.0.24.2) to destination 224.0.0.1 (G073088)
Dec 31 2011 13:18:36.440.1+00:00 R2 MGMD/7/QUERY:Send version 2 general query
on GigabitEthernet0/0/1(10.0.24.2) to destination 224.0.0.1 (G073088)
```

可以看到当健壮系数修改为3以后，当接口启用后前3个普遍组查询消息的时间间隔为5秒，从第四个普遍组查询消息开始时间间隔为20秒。

步骤四.　观察组播路由表

为了观察组播路由的传递，在该拓扑上启用OSPF作为单播路由协议。

```
[R1]ospf 1 router-id 10.0.1.1
[R1-ospf-1]area 0
[R1-ospf-1-area-0.0.0.0]network 10.0.1.1 0.0.0.0
[R1-ospf-1-area-0.0.0.0]network 10.0.14.1 0.0.0.0
[R1-ospf-1-area-0.0.0.0]network 10.0.13.1 0.0.0.0
[R1-ospf-1-area-0.0.0.0]network 10.0.12.1 0.0.0.0

[R2]ospf 1 router-id 10.0.2.2
[R2-ospf-1]area 0
[R2-ospf-1-area-0.0.0.0]network 10.0.2.2 0.0.0.0
[R2-ospf-1-area-0.0.0.0]network 10.0.25.2 0.0.0.0
[R2-ospf-1-area-0.0.0.0]network 10.0.12.2 0.0.0.0

[R3]ospf 1 router-id 10.0.3.3
[R3-ospf-1]area 0
[R3-ospf-1-area-0.0.0.0]network 10.0.3.3 0.0.0.0
[R3-ospf-1-area-0.0.0.0]network 10.0.13.3 0.0.0.0
[R3-ospf-1-area-0.0.0.0]network 10.0.35.3 0.0.0.0

[R4]ospf 1 router-id 10.0.4.4
[R4-ospf-1]area 0
[R4-ospf-1-area-0.0.0.0]network 10.0.4.4 0.0.0.0
[R4-ospf-1-area-0.0.0.0]network 10.0.14.4 0.0.0.0

[R5]ospf 1 router-id 10.0.5.5
[R5-ospf-1]area 0
[R5-ospf-1-area-0.0.0.0]network 10.0.5.5 0.0.0.0
[R5-ospf-1-area-0.0.0.0]network 10.0.25.5 0.0.0.0
[R5-ospf-1-area-0.0.0.0]network 10.0.35.5 0.0.0.0
```

配置完成以后，检查各路由器已能学习到其他路由器的Loopback地址。

```
[R2]display ip routing-table protocol ospf
Route Flags: R - relay, D - download to fib
------------------------------------------------------------------------------
Public routing table : OSPF
        Destinations : 7       Routes : 8

OSPF routing table status : <Active>
        Destinations : 7       Routes : 8

Destination/Mask  Proto Pre  Cost     Flags NextHop       Interface

      10.0.1.1/32  OSPF 10   1562       D   10.0.12.1      Serial1/0/0
      10.0.3.3/32  OSPF 10   1563       D   10.0.12.1      Serial1/0/0
                   OSPF 10   1563       D   10.0.25.5    GigabitEthernet0/0/0
      10.0.4.4/32  OSPF 10   3124       D   10.0.12.1      Serial1/0/0
      10.0.5.5/32  OSPF 10   1          D   10.0.25.5    GigabitEthernet0/0/0
     10.0.13.0/24  OSPF 10   1563       D   10.0.12.1      Serial1/0/0
     10.0.14.0/24  OSPF 10   3124       D   10.0.12.1      Serial1/0/0
     10.0.35.0/24  OSPF 10   1563       D   10.0.25.5    GigabitEthernet0/0/0

OSPF routing table status : <Inactive>
        Destinations : 0       Routes : 0
```

为了模拟组播信息的传递，我们在R3上以自己的Loopback接口作为源地址，向目的地址225.1.1.1发送Ping数据包，模拟组播源。

```
[R3]ping -a 10.0.3.3 -c 300 225.1.1.1
```

间隔几分钟后，我们可以在其他所有路由器上看到组播路由表。在R2上查看组播路由表。

```
[R2]display pim routing-table
 VPN-Instance: public net
 Total 1(*, G) entry; 1 (S, G) entry

(*, 225.1.1.1)
    Protocol: pim-dm, Flag: WC EXT
    UpTime: 00:09:04
    Upstream interface: NULL
        Upstream neighbor: NULL
        RPF prime neighbor: NULL
    Downstream interface(s) information: None

 (10.0.3.3, 225.1.1.1)
```

```
Protocol: pim-dm, Flag:
UpTime: 00:00:52
Upstream interface: GigabitEthernet0/0/0
   Upstream neighbor: 10.0.25.5
   RPF prime neighbor: 10.0.25.5
Downstream interface(s) information: None
```

可以看到2个条目。

第一个条目(*, 225.1.1.1)为该接口配置了静态IGMP组产生的。

第二个条目(10.0.3.3, 225.1.1.1)为组播流量进行扩散后在该路由器上产生的条目。

从输出中我们还可以看到对于R2来说，该组播流的上游路由器为10.0.25.5。

启用了PIM以后，路由器会采用单播路由表进行RPF检查，从下面的输出中可以看到，对于组播源10.0.3.3，RPF的邻居是10.0.25.5。

```
[R2]display multicast rpf-info 10.0.3.3
 VPN-Instance: public net
 RPF information about source: 10.0.3.3
    RPF interface: GigabitEthernet0/0/0, RPF neighbor: 10.0.25.5
    Referenced route/mask: 10.0.3.3/32
    Referenced route type: unicast
    Route selection rule: preference-preferred
    Load splitting rule: disable
```

步骤五. 调整 PIM DM 参数

有时我们希望流量不按照单播路由的路径流向目的地，就可以通过**rpf-route-static**静态修改RPF路径。这个试验中，我们把RPF路径由原来的10.0.25.5修改为10.0.12.1。

```
[R2]ip rpf-route-static 10.0.3.0 255.255.255.0 10.0.12.1
```

配置完成以后，可验证RPF邻居已变成了10.0.12.1。

```
[R2]display multicast rpf-info 10.0.3.3
 VPN-Instance: public net
 RPF information about source: 10.0.3.3
    RPF interface: Serial1/0/0, RPF neighbor: 10.0.12.1
    Referenced route/mask: 10.0.3.0/24
    Referenced route type: mstatic
    Route selection rule: preference-preferred
```

```
    Load splitting rule: disable
```

为了观察PIM的剪枝及嫁接消息，我们通过删除及添加IGMP静态组的方式来模拟用户的离开及加入。首先在R2上打开Debugging。

```
<R2>debugging pim join-prune
<R2>terminal debugging
```

然后把R2的静态IGMP组225.1.1.1删除。

```
[R2-GigabitEthernet0/0/1]undo  igmp static-group 225.1.1.1
Dec 31 2011 15:00:05.300.1+00:00 R2 PIM/7/JP:(public net): PIM ver 2 JP  sending
10.0.12.2 -> 224.0.0.13 on Serial1/0/0  (P012689)
Dec 31 2011 15:00:05.300.2+00:00 R2 PIM/7/JP:(public net): Upstream 10.0.12.1,
Groups 1, Holdtime 210 (P012693)
Dec 31 2011 15:00:05.300.3+00:00 R2 PIM/7/JP:(public net): Group: 225.1.1.1/32
--- 0 joins 1 prunes  (P012701)
Dec 31 2011 15:00:05.310.1+00:00 R2 PIM/7/JP:(public net): Prune: 10.0.3.3/32
(P012707)
Dec 31 2011 15:00:05.350.1+00:00 R2 PIM/7/JP:(public net): PIM ver 2 JP  receiving
10.0.12.1 -> 224.0.0.13 on Serial1/0/0  (P012689)
Dec 31 2011 15:00:05.350.2+00:00 R2 PIM/7/JP:(public net): Upstream 10.0.12.1,
Groups 1, Holdtime 207 (P012693)
Dec 31 2011 15:00:05.350.3+00:00 R2 PIM/7/JP:(public net): Group: 225.1.1.1/32
--- 0 joins 1 prunes (P012701)
Dec 31 2011 15:00:05.350.4+00:00 R2 PIM/7/JP:(public net): Prune: 10.0.3.3/32
(P012707)
```

可以看到R2立刻以组播地址224.0.0.13向上游接口发送剪枝消息，上游路由器的地址为10.0.12.1，此时225.1.1.1这个组播组已被剪枝。随后R1向R2发送消息确认剪枝。

然后再把刚才删除的静态IGMP组播组添加回去。

```
[R2-GigabitEthernet0/0/1] igmp static-group 225.1.1.1
Dec 31 2011 15:00:19.440.1+00:00 R2 PIM/7/JP:(public net): PIM ver 2 GFT sending
10.0.12.2 -> 10.0.12.1 on Serial1/0/0  (P012633)
Dec 31 2011 15:00:19.440.2+00:00 R2 PIM/7/JP:(public net): Upstream 10.0.12.1,
Groups 1, Holdtime 0 (P012639)
Dec 31 2011 15:00:19.440.3+00:00 R2 PIM/7/JP:(public net): Group: 225.1.1.1/32
--- 1 joins 0 prunes  (P012648)
Dec 31 2011 15:00:19.440.4+00:00 R2 PIM/7/JP:(public net): Join: 10.0.3.3/32
(P012654)
Dec 31 2011 15:00:19.480.1+00:00 R2 PIM/7/JP:(public net): PIM ver 2 GAK receiving
10.0.12.1 -> 10.0.12.2 on Serial1/0/0  (P012633)
```

```
Dec 31 2011 15:00:19.480.2+00:00 R2 PIM/7/JP:(public net): Upstream 10.0.12.2,
Groups 1, Holdtime 0 (P012639)
Dec 31 2011 15:00:19.480.3+00:00 R2 PIM/7/JP:(public net): Group: 225.1.1.1/32
--- 1 joins 0 prunes (P012648)
Dec 31 2011 15:00:19.480.4+00:00 R2 PIM/7/JP:(public net): Join: 10.0.3.3/32
(P012654)
```

这时R2立刻向上游以单播的形式发送了嫁接消息，加入225.1.1.1，同时R1也以单播的形式向R2回应了嫁接确认。

从这里可以总结出：**剪枝消息是以组播地址224.0.0.13发送的，而嫁接消息是以单播向上游发送的。**

有时我们希望组播流量只在规定的范围内传递，这时候可以在接口下通过**multicast boundary**为某个特定的组播组或组播地址段定界。

控制组播组225.1.1.2的流量不要传递到R4上，在R1连接到R4的接口上增加如下配置。

```
[R1-Serial3/0/0]multicast boundary 225.1.1.2 255.255.255.255
```

在R3上模拟目的地址为225.1.1.2的组播流量。

```
[R3]ping -a 10.0.3.3 -c 300 225.1.1.2
```

等待在R2和R4上分别查看组播路由表，可以看到在R2上存在表项(10.0.3.3, 225.1.1.2)，而在R4上没有该组播组路由条目，说明组播流量并没有扩散到R4上。

```
[R2]display pim routing-table
 VPN-Instance: public net
 Total 1 (*, G) entry; 2 (S, G) entries

 (*, 225.1.1.1)
    Protocol: pim-dm, Flag: WC EXT
    UpTime: 00:09:04
    Upstream interface: NULL
        Upstream neighbor: NULL
        RPF prime neighbor: NULL
    Downstream interface(s) information: None

 (10.0.3.3, 225.1.1.1)
    Protocol: pim-dm, Flag: EXT
    UpTime: 00:02:11
    Upstream interface: Serial1/0/0
        Upstream neighbor: 10.0.12.1
```

```
      RPF prime neighbor: 10.0.12.1
   Downstream interface(s) information: None

 (10.0.3.3, 225.1.1.2)
   Protocol: pim-dm, Flag:
   UpTime: 00:00:08
   Upstream interface: Serial1/0/0
      Upstream neighbor: 10.0.12.1
      RPF prime neighbor: 10.0.12.1
   Downstream interface(s) information: None

[R4]display pim routing-table
 VPN-Instance: public net
 Total 1 (*, G) entry; 1 (S, G) entry

 (*, 225.1.1.1)
   Protocol: pim-dm, Flag: WC
   UpTime: 00:08:03
   Upstream interface: NULL
      Upstream neighbor: NULL
      RPF prime neighbor: NULL
   Downstream interface(s) information:
   Total number of downstreams: 1
      1: GigabitEthernet0/0/1
         Protocol: static, UpTime: 00:08:03, Expires: never

 (10.0.3.3, 225.1.1.1)
   Protocol: pim-dm, Flag:
   UpTime: 00:02:43
   Upstream interface: Serial1/0/0
      Upstream neighbor: 10.0.14.1
      RPF prime neighbor: 10.0.14.1
   Downstream interface(s) information:
   Total number of downstreams: 1
      1: GigabitEthernet0/0/1
         Protocol: pim-dm, UpTime: 00:02:43, Expires:  -
```

默认情况下PIM DM选取接口IP地址较大的路由器作为DR。

```
[R2]display pim interface
 VPN-Instance: public net
 Interface          State NbrCnt HelloInt  DR-Pri   DR-Address
 GE0/0/0              up    1      30         1         10.0.25.5
 S1/0/0               up    1      30         1         10.0.12.2       (local)
```

在R2上查看接口状态可以看到在与R5连接的接口上，R5是DR。我们可以通过修改接口的优先级来影响DR的选举，该优先级值是一个32bit长度的数值，默认值为1。在下面的例子中，将R2连接到R5的接口的优先级改成100。

```
[R2-GigabitEthernet0/0/0]pim hello-option dr-priority 100

[R2]display pim interface
 VPN-Instance: public net
 Interface           State NbrCnt HelloInt DR-Pri     DR-Address
 GE0/0/0              up    1      30        100        10.0.25.2      (local)
 S1/0/0               up    1      30        1          10.0.12.2      (local)
```

可以看到当把路由器接口优先级更改到100以后，R2立刻抢占了DR的位置。

有时为了安全需要，我们希望面向用户侧的接口上不再收发PIM的Hello包，使用**pim silent**可实现该功能。

```
[R4-GigabitEthernet0/0/1]pim silent
```

配置完成以后检查PIM Silent已生效。

```
[R4]display pim interface GigabitEthernet 0/0/1 verbose
 VPN-Instance: public net
 Interface: GigabitEthernet0/0/1, 10.0.24.4
     PIM version: 2
     PIM mode: Dense
     PIM state: up
     PIM DR: 10.0.24.4 (local)
     PIM DR Priority (configured): 1
     PIM neighbor count: 0
     PIM hello interval: 30 s
     PIM LAN delay (negotiated): 500 ms
     PIM LAN delay (configured): 500 ms
     PIM hello override interval (negotiated): 2500 ms
     PIM hello override interval (configured): 2500 ms
     PIM Silent: enabled
     PIM neighbor tracking (negotiated): disabled
     PIM neighbor tracking (configured): disabled
     PIM generation ID: 0XAD457D14
     PIM require-GenID: disabled
     PIM hello hold interval: 105 s
     PIM assert hold interval: 180 s
     PIM triggered hello delay: 5 s
     PIM J/P interval: 60 s
```

```
    PIM J/P hold interval: 210 s
    PIM state-refresh processing: enabled
    PIM state-refresh interval: 60 s
    PIM graft retry interval: 3 s
    PIM state-refresh capability on link: capable
    PIM dr-switch-delay timer : not configured
    Number of routers on link not using DR priority: 0
    Number of routers on link not using LAN delay: 0
    Number of routers on link not using neighbor tracking: 1
    ACL of PIM neighbor policy: -
    ACL of PIM ASM join policy: -
    ACL of PIM SSM join policy: -
    ACL of PIM join policy: -
```

附加实验：思考并验证

PIM的DM模式适合于用户比较多，比较密集的场景。

思考一下生活中哪些网络应用适合使用PIM DM模式的组播来实现数据转发？它们的特点是什么？

PIM的DM模式在应用在大规模网络上，有哪些劣势？

最终设备配置

```
<R1>display current-configuration
[V200R001C00SPC200]
#
 sysname R1
#
interface Serial1/0/0
 link-protocol ppp
 ip address 10.0.12.1 255.255.255.0
 pim dm
#
interface Serial3/0/0
 link-protocol ppp
 ip address 10.0.14.1 255.255.255.0
 pim dm
```

```
 multicast boundary 225.1.1.2 32
#
ip address 10.0.13.1 255.255.255.0
 pim dm
#
interface LoopBack0
 ip address 10.0.1.1 255.255.255.255
#
ospf 1 router-id 10.0.1.1
 area 0.0.0.0
  network 10.0.1.1 0.0.0.0
  network 10.0.14.1 0.0.0.0
  network 10.0.13.1 0.0.0.0
  network 10.0.12.1 0.0.0.0
#
return

<R2>display current-configuration
[V200R001C00SPC200]
#
 sysname R2
#
interface Serial1/0/0
 link-protocol ppp
 ip address 10.0.12.2 255.255.255.0
 pim dm
#
interface GigabitEthernet0/0/0
 ip address 10.0.25.2 255.255.255.0
 pim hello-option dr-priority 100
 pim dm
#
interface GigabitEthernet0/0/1
 ip address 10.0.24.2 255.255.255.0
 igmp enable
 igmp robust-count 3
 igmp timer query 20
 igmp static-group 225.1.1.1
#
interface LoopBack0
 ip address 10.0.2.2 255.255.255.255
#
ospf 1 router-id 10.0.2.2
```

```
 area 0.0.0.0
  network 10.0.2.2 0.0.0.0
  network 10.0.25.2 0.0.0.0
  network 10.0.12.2 0.0.0.0
#
 ip rpf-route-static 10.0.3.0 24 10.0.12.1
#
return
```

```
<R3>display current-configuration
[V200R001C00SPC200]
#
 sysname R3
#
interface Serial3/0/0
 link-protocol ppp
 ip address 10.0.35.3 255.255.255.0
 pim dm
#
interface GigabitEthernet0/0/2
 ip address 10.0.13.3 255.255.255.0
 pim dm
#
interface LoopBack0
 ip address 10.0.3.3 255.255.255.255
#
ospf 1 router-id 10.0.3.3
 area 0.0.0.0
  network 10.0.3.3 0.0.0.0
  network 10.0.13.3 0.0.0.0
  network 10.0.35.3 0.0.0.0
#
return
```

```
<R4>display current-configuration
[V200R001C00SPC200]
#
 sysname R4
#
interface Serial1/0/0
 link-protocol ppp
 ip address 10.0.14.4 255.255.255.0
 pim dm
```

```
#
interface GigabitEthernet0/0/1
 ip address 10.0.24.4 255.255.255.0
 pim silent
 igmp enable
 igmp static-group 225.1.1.1
#
interface LoopBack0
 ip address 10.0.4.4 255.255.255.255
#
ospf 1 router-id 10.0.4.4
 area 0.0.0.0
  network 10.0.4.4 0.0.0.0
  network 10.0.14.4 0.0.0.0
#
return

<R5>display current-configuration
[V200R001C00SPC200]
#
 sysname R5
#
interface Serial1/0/0
 link-protocol ppp
 ip address 10.0.35.5 255.255.255.0
 pim dm
#
interface GigabitEthernet0/0/0
 ip address 10.0.25.5 255.255.255.0
 pim dm
#
interface LoopBack0
 ip address 10.0.5.5 255.255.255.255
#
ospf 1 router-id 10.0.5.5
 area 0.0.0.0
  network 10.0.5.5 0.0.0.0
  network 10.0.25.5 0.0.0.0
  network 10.0.35.5 0.0.0.0
#
return
```

实验 5-2 PIM SM 及动态 RP

学习目的

- 掌握PIM SM协议的配置方法
- 掌握静态RP和RP负载均衡的配置方法
- 掌握RPT与SPT之间切换的控制方法
- 掌握Auto-RP的配置方法

拓扑图

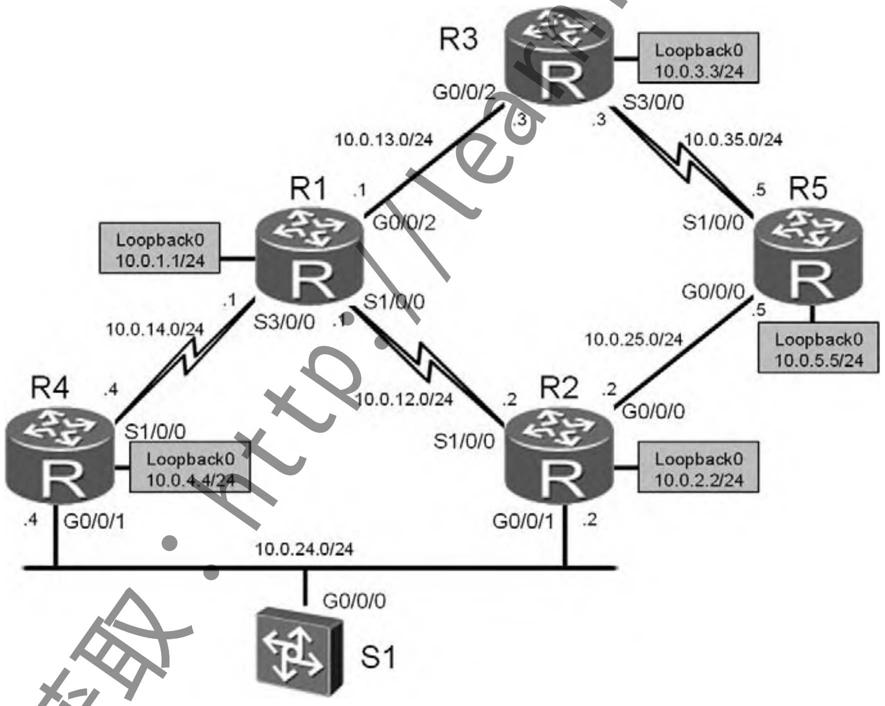

图5-2 PIM SM及动态RP

场景

你是公司的网络管理员。公司以前使用过PIM的DM模式来实现组播路由学习，但是后来发现随着组播应用的发展，组播用户分部越来越广，这时组播质量显示出一定程度的下降。为了提高组播的可靠性、安全性和效率，你决定使用PIM的SM模式来实现组播路由学习。

在PIM的SM模式中，你需要定义RP，作为SM模式的共享树的树根。但是在实际的使用中，只是这样简单的配置是不够的，组播还需要实现RP之间的负载分担。并且考虑到以后组播规模的增大，你需要考虑使用Auto-RP功能。

在实现网络的组播转发之前，你也遇到了一些网络故障，经过一些故障排除步骤，最终网络正常工作了。

学习任务

步骤一. 基础配置与 IP 编址

S2参与到本次实验，但无需配置。实验之前，请清空S2的配置，并重启它。

给所有路由器配置IP地址和掩码。配置时注意所有的Loopback接口配置掩码均为32位。

```
<Huawei>system-view
Enter system view, return user view with Ctrl+Z.
[Huawei]sysname R1
[R1]interface GigabitEthernet0/0/2
[R1-GigabitEthernet0/0/2]ip address 10.0.13.1 24
[R1-GigabitEthernet0/0/2]interface Serial 1/0/0
[R1-Serial1/0/0]ip address 10.0.12.1 24
[R1-Serial1/0/0]interface Serial 3/0/0
[R1-Serial3/0/0]ip address 10.0.14.1 24
[R1-Serial3/0/0]interface loopback 0
[R1-LoopBack0]ip address 10.0.1.1 32

<Huawei>system-view
Enter system view, return user view with Ctrl+Z.
[Huawei]sysname R2
[R2]interface GigabitEthernet0/0/0
[R2-GigabitEthernet0/0/0]ip address 10.0.25.2 24
[R2-GigabitEthernet0/0/0]interface GigabitEthernet0/0/1
[R2-GigabitEthernet0/0/1]ip address 10.0.24.2 24
[R2-GigabitEthernet0/0/1]interface Serial 1/0/0
```

```
[R2-Serial1/0/0]ip address 10.0.12.2 24
[R2-Serial1/0/0]interface loopback 0
[R2-LoopBack0]ip address 10.0.2.2 32

<Huawei>system-view
Enter system view, return user view with Ctrl+Z.
[Huawei]sysname R3
[R3]interface GigabitEthernet0/0/2
[R3-GigabitEthernet0/0/2]ip address 10.0.13.3 24
[R3-GigabitEthernet0/0/2]interface Serial 3/0/0
[R3-Serial3/0/0]ip address 10.0.35.3 24
[R3-Serial3/0/0]interface loopback 0
[R3-LoopBack0]ip address 10.0.3.3 32

<Huawei>system-view
Enter system view, return user view with Ctrl+Z.
[Huawei]sysname R4
[R4]interface GigabitEthernet0/0/1
[R4-GigabitEthernet0/0/1]ip address 10.0.24.4 24
[R4-GigabitEthernet0/0/1]interface Serial 1/0/0
[R4-Serial1/0/0]ip address 10.0.14.4 24
[R4-Serial1/0/0]interface loopback 0
[R4-LoopBack0]ip address 10.0.4.4 32

<Huawei>system-view
Enter system view, return user view with Ctrl+Z.
[Huawei]sysname R5
[R5]interface GigabitEthernet0/0/0
[R5-GigabitEthernet0/0/0]ip address 10.0.25.5 24
[R5-GigabitEthernet0/0/0]interface Serial 1/0/0
[R5-Serial1/0/0]ip add 10.0.35.5 24
[R5-Serial1/0/0]interface loopback 0
[R5-LoopBack0]ip add 10.0.5.5 32

<Quidway>system-view
Enter system view, return user view with Ctrl+Z.
[Quidway]sysname S1
[S1]interface Vlanif 1
[S1-Vlanif1]ip address 10.0.24.1 24
[S1-Vlanif1]interface loopback 0
[S1-LoopBack0]ip address 10.0.11.11 24
```

配置完成后，测试直连链路的连通性。

```
[R1]ping -c 1 10.0.12.2
  PING 10.0.12.2: 56  data bytes, press CTRL_C to break
    Reply from 10.0.12.2: bytes=56 Sequence=1 ttl=255 time=41 ms

  --- 10.0.12.2 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 41/41/41 ms

[R1]ping -c 1 10.0.13.3
  PING 10.0.13.3: 56  data bytes, press CTRL_C to break
    Reply from 10.0.13.3: bytes=56 Sequence=1 ttl=255 time=5 ms

  --- 10.0.13.3 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 5/5/5 ms

[R1]ping -c 1 10.0.14.4
  PING 10.0.14.4: 56  data bytes, press CTRL_C to break
    Reply from 10.0.14.4: bytes=56 Sequence=1 ttl=255 time=62 ms

  --- 10.0.14.4 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 62/62/62 ms

[R5]ping -c 1 10.0.25.2
  PING 10.0.25.2: 56 data bytes, press CTRL_C to break
    Reply from 10.0.25.2: bytes=56 Sequence=1 ttl=255 time=7 ms

  --- 10.0.25.2 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 7/7/7 ms

[R5]ping -c 1 10.0.35.3
  PING 10.0.35.3: 56 data bytes, press CTRL_C to break
    Reply from 10.0.35.3: bytes=56 Sequence=1 ttl=255 time=37 ms
```

```
  --- 10.0.35.3 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 37/37/37 ms

[S1]ping -c 1 10.0.24.2
  PING 10.0.24.2: 56  data bytes, press CTRL_C to break
    Reply from 10.0.24.2: bytes=56 Sequence=1 ttl=255 time=1 ms

  --- 10.0.24.2 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 1/1/1 ms
```

在R1、R2、R3、R4、 R5和S1上启用OSPF路由协议。实现所有网络互通。

```
[R1]ospf 1
[R1-ospf-1]area 0
[R1-ospf-1-area-0.0.0.0]network 10.0.14.1 0.0.0.0
[R1-ospf-1-area-0.0.0.0]network 10.0.12.1 0.0.0.0
[R1-ospf-1-area-0.0.0.0]network 10.0.13.1 0.0.0.0
[R1-ospf-1-area-0.0.0.0]network 10.0.1.1 0.0.0.0

[R2]ospf 1
[R2-ospf-1]area 0
[R2-ospf-1-area-0.0.0.0]network 10.0.12.2 0.0.0.0
[R2-ospf-1-area-0.0.0.0]network 10.0.24.2 0.0.0.0
[R2-ospf-1-area-0.0.0.0]network 10.0.25.2 0.0.0.0
[R2-ospf-1-area-0.0.0.0]network 10.0.2.2 0.0.0.0

[R3]ospf 1
[R3-ospf-1]area 0
[R3-ospf-1-area-0.0.0.0]network 10.0.13.3 0.0.0.0
[R3-ospf-1-area-0.0.0.0]network 10.0.35.3 0.0.0.0
[R3-ospf-1-area-0.0.0.0]network 10.0.3.3 0.0.0.0

[R4]ospf 1
[R4-ospf-1]area 0
[R4-ospf-1-area-0.0.0.0]network 10.0.14.4 0.0.0.0
[R4-ospf-1-area-0.0.0.0]network 10.0.24.4 0.0.0.0
[R4-ospf-1-area-0.0.0.0]network 10.0.4.4 0.0.0.0
```

```
[R5]ospf 1
[R5-ospf-1]area 0
[R5-ospf-1-area-0.0.0.0]network 10.0.25.5 0.0.0.0
[R5-ospf-1-area-0.0.0.0]network 10.0.35.5 0.0.0.0
[R5-ospf-1-area-0.0.0.0]network 10.0.5.5 0.0.0.0

[S1]ospf 1
[S1-ospf-1]area 0
[S1-ospf-1-area-0.0.0.0]network 10.0.24.1 0.0.0.0
[S1-ospf-1-area-0.0.0.0]network 10.0.11.11 0.0.0.0
```

配置完成后，稍等片刻，待OSPF邻居关系连接，路由信息交互完成后，测试S1与路由器R3的Loopback 0地址之间的连通性。

```
[S1]ping -c 1 10.0.3.3
  PING 10.0.3.3: 56  data bytes, press CTRL_C to break
    Reply from 10.0.3.3: bytes=56 Sequence=1 ttl=253 time=37 ms

  --- 10.0.3.3 ping statistics ---
    1 packet(s) transmitted
    1 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 37/37/37 ms
```

如测试显示，网络工作正常。

步骤二. 配置所有路由器启用 PIM SM

启用R1、R2、R3、R4、R5和S1的组播路由功能。

```
[R1]multicast routing-enable

[R2]multicast routing-enable

[R3]multicast routing-enable

[R4]multicast routing-enable

[R5]multicast routing-enable

[S1]multicast routing-enable
```

在所有设备的所有接口上配置运行SM模式的PIM。

```
[R1]interface GigabitEthernet0/0/2
[R1-GigabitEthernet0/0/2]pim sm
[R1-GigabitEthernet0/0/2]interface Serial 1/0/0
[R1-Serial1/0/0]pim sm
[R1-Serial1/0/0]interface Serial 3/0/0
[R1-Serial3/0/0]pim sm
[R1-Serial3/0/0]interface loopback 0
[R1-LoopBack0]pim sm

[R2]interface GigabitEthernet0/0/0
[R2-GigabitEthernet0/0/0]pim sm
[R2-GigabitEthernet0/0/0]interface GigabitEthernet0/0/1
[R2-GigabitEthernet0/0/1]pim sm
[R2-GigabitEthernet0/0/1]interface Serial 1/0/0
[R2-Serial1/0/0]pim sm
[R2-Serial1/0/0]interface loopback 0
[R2-LoopBack0]pim sm

[R3]interface GigabitEthernet0/0/2
[R3-GigabitEthernet0/0/2]pim sm
[R3-GigabitEthernet0/0/2]interface Serial 3/0/0
[R3-Serial3/0/0]pim sm
[R3-Serial3/0/0]interface loopback 0
[R3-LoopBack0]pim sm

[R4]interface GigabitEthernet0/0/1
[R4-GigabitEthernet0/0/1]pim sm
[R4-GigabitEthernet0/0/1]interface Serial 1/0/0
[R4-Serial1/0/0]pim sm
[R4-Serial1/0/0]interface loopback 0
[R4-LoopBack0]pim sm

[R5]interface GigabitEthernet0/0/0
[R5-GigabitEthernet0/0/0]pim sm
[R5-GigabitEthernet0/0/0]interface Serial 1/0/0
[R5-Serial1/0/0]pim sm
[R5-Serial1/0/0]interface loopback 0
[R5-LoopBack0]pim sm

[S1]interface Vlanif 1
[S1-Vlanif1]pim sm
[S1-Vlanif1]interface loopback 0
```

```
[S1-LoopBack0]pim sm
```

配置完成后，查看R1、R5和S1的PIM邻居学习的情况。

```
<R1>display pim neighbor
 VPN-Instance: public net
 Total Number of Neighbors = 3

 Neighbor        Interface          Uptime   Expires  Dr-Priority
 10.0.13.3       GE0/0/2            00:08:52 00:01:23 1
 10.0.12.2       S1/0/0             00:40:44 00:01:30 1
 10.0.14.4       S3/0/0             00:07:53 00:01:23 1

[R5]display pim neighbor
 VPN-Instance: public net
 Total Number of Neighbors = 2

 Neighbor        Interface          Uptime   Expires  Dr-Priority
 10.0.25.2       GE0/0/0            00:08:38 00:01:30 1
 10.0.35.3       S1/0/0             00:08:38 00:01:28 1

[S1]display pim neighbor
 VPN-Instance: public net
 Total Number of Neighbors = 2

 Neighbor        Interface          Uptime   Expires  Dr-Priority  BFD-Session
 10.0.24.4       Vlanif1            00:01:24 00:01:23 1            N
 10.0.24.2       Vlanif1            00:01:22 00:01:17 1            N
```

从命令输出结果可以看到，PIM协议已经在网络中正确运行。

步骤三.　静态 RP 和静态 RP 负载均衡

通过给网络手动指定静态RP来控制网络中组播数据流。

在所有设备配置R1的S3/0/0接口作为网络中的静态RP。

```
[R1]pim
[R1-pim]static-rp 10.0.14.1

[R2]pim
[R2-pim]static-rp 10.0.14.1

[R3]pim
```

```
[R3-pim]static-rp 10.0.14.1

[R4]pim
[R4-pim]static-rp 10.0.14.1

[R5]pim
[R5-pim]static-rp 10.0.14.1

[S1]pim
[S1-pim]static-rp 10.0.14.1
```

S1模拟网络中组播用户接入的三层交换机，在S1的Loopback 0接口上开启IGMP功能。

```
[S1]interface LoopBack 0
[S1-LoopBack0]igmp enable
```

将S1的Loopback 0接口静态加入225.0.0.1组播组，模拟连接有225.0.0.1组播组的用户。

```
[S1]interface LoopBack 0
[S1-LoopBack0]igmp static-group 225.0.0.1
```

在R1、R4和S1上使用命令**display pim routing-table**查看PIM路由表。

```
[R1]display pim routing-table
 VPN-Instance: public net
 Total 1 (*, G) entry; 0 (S, G) entry

 (*, 225.0.0.1)
     RP: 10.0.14.1 (local)
     Protocol: pim-sm, Flag: WC
     UpTime: 00:02:40
     Upstream interface: Register
         Upstream neighbor: NULL
         RPF prime neighbor: NULL
     Downstream interface(s) information:
     Total number of downstreams: 1
         1: Serial3/0/0
             Protocol: pim-sm, UpTime: 00:02:40, Expires: 00:02:50

[R4]display pim routing-table
 VPN-Instance: public net
 Total 1 (*, G) entry; 0 (S, G) entry
```

```
 (*, 225.0.0.1)
    RP: 10.0.14.1
    Protocol: pim-sm, Flag: WC
    UpTime: 00:01:46
    Upstream interface: Serial1/0/0
        Upstream neighbor: 10.0.14.1
        RPF prime neighbor: 10.0.14.1
    Downstream interface(s) information:
    Total number of downstreams: 1
        1: GigabitEthernet0/0/1
            Protocol: pim-sm, UpTime: 00:01:46, Expires: 00:02:43

[S1-LoopBack0]display pim routing-table
 VPN-Instance: public net
 Total 1 (*, G) entry; 0 (S, G) entry

 (*, 225.0.0.1)
    RP: 10.0.14.1
    Protocol: pim-sm, Flag: WC
    UpTime: 00:01:19
    Upstream interface: Vlanif1
        Upstream neighbor: 10.0.24.4
        RPF prime neighbor: 10.0.24.4
    Downstream interface(s) information:
    Total number of downstreams: 1
        1: LoopBack0
            Protocol: static, UpTime: 00:01:19, Expires: -
```

从命令输出结果可以看到，在指定静态RP的网络中，R1是网络中的RP。S1生成了一条经过R4到达RP路由器R1的组播路径。

创建ACL并应用到静态RP上，定义R1作为RP，服务的组播范围是225.0.0.0/24网段，定义R5作为RP，服务的组播范围是225.0.1.0/24网段。

```
[R1]acl 2000
[R1-acl-basic-2000]rule permit source 225.0.0.0 0.0.0.255
[R1-acl-basic-2000]acl 2001
[R1-acl-basic-2001]rule permit source 225.0.1.0 0.0.0.255
[R1-acl-basic-2001]pim
[R1-pim]static-rp 10.0.14.1 2000
[R1-pim]static-rp 10.0.25.5 2001

[R2]acl 2000
[R2-acl-basic-2000]rule permit source 225.0.0.0 0.0.0.255
```

```
[R2-acl-basic-2000]acl 2001
 [R2-acl-basic-2001]rule permit source 225.0.1.0 0.0.0.255
[R2-acl-basic-2001]pim
 [R2-pim]static-rp 10.0.14.1 2000
[R2-pim]static-rp 10.0.25.5 2001

[R3]acl 2000
[R3-acl-basic-2000]rule permit source 225.0.0.0 0.0.0.255
[R3-acl-basic-2000]acl 2001
 [R3-acl-basic-2001]rule permit source 225.0.1.0 0.0.0.255
[R3-acl-basic-2001]pim
[R3-pim]static-rp 10.0.14.1 2000
[R3-pim]static-rp 10.0.25.5 2001

[R4]acl 2000
[R4-acl-basic-2000]rule permit source 225.0.0.0 0.0.0.255
[R4-acl-basic-2000]acl 2001
[R4-acl-basic-2001]rule permit source 225.0.1.0 0.0.0.255
[R4-acl-basic-2001]pim
[R4-pim]static-rp 10.0.14.1 2000
[R4-pim]static-rp 10.0.25.5 2001

[R5]acl 2000
[R5-acl-basic-2000]rule permit source 225.0.0.0 0.0.0.255
[R5-acl-basic-2000]acl 2001
[R5-acl-basic-2001]rule permit source 225.0.1.0 0.0.0.255
[R5-acl-basic-2001]pim
[R5-pim]static-rp 10.0.14.1 2000
[R5-pim]static-rp 10.0.25.5 2001

[S1]acl 2000
[S1-acl-basic-2000]rule permit source 225.0.0.0 0.0.0.255
[S1-acl-basic-2000]acl 2001
[S1-acl-basic-2001]rule permit source 225.0.1.0 0.0.0.255
[S1-acl-basic-2001]pim
[S1-pim]static-rp 10.0.14.1 2000
[S1-pim]static-rp 10.0.25.5 2001
```

将S1的Loopback 0接口静态加入225.0.1.1组播组，模拟连接有225.0.1.1组播组的用户。

```
[S1]interface LoopBack 0
[S1-LoopBack0]igmp static-group 225.0.1.1
```

在S1、R2和 R5上使用命令**display pim routing-table**查看PIM路由表。

```
[R5]dis pim routing-table
 VPN-Instance: public net
 Total 1 (*, G) entry; 0 (S, G) entry

 (*, 225.0.1.1)
     RP: 10.0.25.5 (local)
     Protocol: pim-sm, Flag: WC
     UpTime: 00:03:13
     Upstream interface: Register
         Upstream neighbor: NULL
         RPF prime neighbor: NULL
     Downstream interface(s) information:
     Total number of downstreams: 1
         1: GigabitEthernet0/0/0
             Protocol: pim-sm, UpTime: 00:03:13, Expires: 00:03:17

[R2]display pim routing-table
 VPN-Instance: public net
 Total 1 (*, G) entry; 0 (S, G) entry

 (*, 225.0.1.1)
     RP: 10.0.25.5
     Protocol: pim-sm, Flag: WC
     UpTime: 00:03:41
     Upstream interface: GigabitEthernet0/0/0
         Upstream neighbor: 10.0.25.5
         RPF prime neighbor: 10.0.25.5
     Downstream interface(s) information:
     Total number of downstreams: 1
         1: GigabitEthernet0/0/1
             Protocol: pim-sm, UpTime: 00:03:41, Expires: 00:02:48

[S1]display pim routing-table
 VPN-Instance: public net
 Total 2 (*, G) entries; 0 (S, G) entry

 (*, 225.0.0.1)
     RP: 10.0.14.1
     Protocol: pim-sm, Flag: WC
     UpTime: 00:17:09
     Upstream interface: Vlanif1
         Upstream neighbor: 10.0.24.4
```

```
        RPF prime neighbor: 10.0.24.4
     Downstream interface(s) information:
     Total number of downstreams: 1
        1: LoopBack0
           Protocol: static, UpTime: 00:17:09, Expires: -

 (*, 225.0.1.1)
     RP: 10.0.25.5
     Protocol: pim-sm, Flag: WC
     UpTime: 00:03:58
     Upstream interface: Vlanif1
        Upstream neighbor: 10.0.24.2
        RPF prime neighbor: 10.0.24.2
     Downstream interface(s) information:
     Total number of downstreams: 1
        1: LoopBack0
           Protocol: static, UpTime: 00:03:58, Expires: -
```

从命令输出结果可以看到，S1针对225.0.0.1和225.0.1.1生成了二条组播路径。225.0.1.1组播路径经过R2到达RP路由器R5。

步骤四. 配置 Auto-RP

使用Auto-RP方式配置R3作为C-BSR控制组播边界，同时让PIM自动在R1和R5之间选举出RP。

删除所有路由器上的静态RP配置，避免静态配置干扰Auto-RP实验。

```
[R1]undo pim
Warning: This operation will lead to the deletion of all the IPv4 global PIM
configurations in the public instance. Continue? [Y/N]:y

[R2]undo pim
Warning: This operation will lead to the deletion of all the IPv4 global PIM
configurations in the public instance. Continue? [Y/N]:y

[R3]undo pim
Warning: This operation will lead to the deletion of all the IPv4 global PIM
configurations in the public instance. Continue? [Y/N]:y

[R4]undo pim
Warning: This operation will lead to the deletion of all the IPv4 global PIM
configurations in the public instance. Continue? [Y/N]:y
```

```
[R5]undo pim
Warning: This operation will lead to the deletion of all the IPv4 global PIM
configurations in the public instance. Continue? [Y/N]:y

[S1]undo pim
```

在R1和R5上配置Auto-RP。

```
[R1]pim
[R1-pim]c-rp LoopBack 0

[R5]pim
[R5-pim]c-rp LoopBack 0
```

将R3配置为C-BSR。

```
[R3]pim
[R3-pim]c-bsr LoopBack 0
```

在R1上使用**display pim bsr-info**命令查看网络中的C-BSR信息。

```
[R1]display pim bsr-info
 VPN-Instance: public net
 Elected AdminScoped BSR Count: 0
 Elected BSR Address: 10.0.3.3
     Priority: 0
     Hash mask length: 30
     State: Accept Preferred
     Scope: Not scoped
     Uptime: 00:02:46
     Expires: 00:01:34
     C-RP Count: 2
```

在R1上使用display pim rp-info命令查看网络中的RP信息。

```
[R1]display pim rp-info
 VPN-Instance: public net
 PIM-SM BSR RP Number:2
 Group/MaskLen: 224.0.0.0/4
     RP: 10.0.1.1 (local)
     Priority: 0
     Uptime: 00:04:51
     Expires: 00:01:39
 Group/MaskLen: 224.0.0.0/4
     RP: 10.0.5.5
```

```
    Priority: 0
    Uptime: 00:04:51
    Expires: 00:01:39
```

在S1上查看PIM信息。

```
[S1]display pim routing-table
 VPN-Instance: public net
 Total 2 (*, G) entries; 0 (S, G) entry

 (*, 225.0.0.1)
    RP: 10.0.5.5
    Protocol: pim-sm, Flag: WC
    UpTime: 00:31:32
    Upstream interface: Vlanif1
       Upstream neighbor: 10.0.24.2
       RPF prime neighbor: 10.0.24.2
    Downstream interface(s) information:
    Total number of downstreams: 1
       1: LoopBack0
          Protocol: static, UpTime: 00:31:32, Expires: -

 (*, 225.0.1.1)
    RP: 10.0.1.1
    Protocol: pim-sm, Flag: WC
    UpTime: 00:18:21
    Upstream interface: Vlanif1
       Upstream neighbor: 10.0.24.4
       RPF prime neighbor: 10.0.24.4
    Downstream interface(s) information:
    Total number of downstreams: 1
       1: LoopBack0
          Protocol: static, UpTime: 00:18:21, Expires: -
```

从输出结果可以看到，R3是网络中的C-BSR，R1是225.0.1.1组播组的RP，R5是225.0.0.1组播组的RP。225.0.1.1组播组的路径是从R4到R1，225.0.0.1组播组的路径是从R2到R5。

附加实验：思考并验证

PIM的SM模式适合于用户比较分散的场景。

思考一下生活中哪些网络应用适合使用PIM SM模式的组播来实现数据转发？

它们的特点是什么？

最终设备配置

```
[R1]display current-configuration
[V200R001C00SPC200]
#
 sysname R1
#
 board add 0/1 1SA
 board add 0/2 1SA
 board add 0/3 1SA
#
 multicast routing-enable
#
acl number 2000
 rule 5 permit source 225.0.0.0 0.0.0.255
#
acl number 2001
 rule 5 permit source 225.0.1.0 0.0.0.255
#
interface Serial1/0/0
 link-protocol ppp
 ip address 10.0.12.1 255.255.255.0
 pim sm
#
interface Serial3/0/0
 link-protocol ppp
 ip address 10.0.14.1 255.255.255.0
 pim sm
#
interface GigabitEthernet0/0/2
 ip address 10.0.13.1 255.255.255.0
 pim sm
#
interface LoopBack0
 ip address 10.0.1.1 255.255.255.255
 pim sm
#
ospf 1
 area 0.0.0.0
  network 10.0.14.1 0.0.0.0
  network 10.0.12.1 0.0.0.0
```

```
  network 10.0.13.1 0.0.0.0
  network 10.0.1.1 0.0.0.0
#
pim
 c-rp LoopBack0
#
Return
```

```
[R2]display current-configuration
[V200R001C00SPC200]
#
 sysname R2
#
 board add 0/1 1SA
 board add 0/2 1SA
 board add 0/3 1SA
#
 multicast routing-enable
#
acl number 2000
 rule 5 permit source 225.0.0.0 0.0.0.255
#
acl number 2001
 rule 5 permit source 225.0.1.0 0.0.0.255
#
interface Serial1/0/0
 link-protocol ppp
 ip address 10.0.12.2 255.255.255.0
 pim sm
#
interface GigabitEthernet0/0/0
 ip address 10.0.25.2 255.255.255.0
 pim sm
#
interface GigabitEthernet0/0/1
 ip address 10.0.24.2 255.255.255.0
 pim sm
#
interface LoopBack0
 ip address 10.0.2.2 255.255.255.255
 pim sm
#
ospf 1
```

```
 area 0.0.0.0
  network 10.0.12.2 0.0.0.0
  network 10.0.24.2 0.0.0.0
  network 10.0.25.2 0.0.0.0
  network 10.0.2.2 0.0.0.0
#
Return
```

```
[R3]display current-configuration
[V200R001C00SPC200]
#
 sysname R3
#
 board add 0/1 1SA
 board add 0/2 1SA
 board add 0/3 1SA
#
 multicast routing-enable
#
acl number 2000
 rule 5 permit source 225.0.0.0 0.0.0.255
#
acl number 2001
 rule 5 permit source 225.0.1.0 0.0.0.255
#
interface Serial3/0/0
 link-protocol ppp
 ip address 10.0.35.3 255.255.255.0
 pim sm
#
interface GigabitEthernet0/0/2
 ip address 10.0.13.3 255.255.255.0
 pim sm
#
interface LoopBack0
 ip address 10.0.3.3 255.255.255.255
 pim sm
#
ospf 1
 area 0.0.0.0
  network 10.0.13.3 0.0.0.0
  network 10.0.35.3 0.0.0.0
  network 10.0.3.3 0.0.0.0
```

```
#
pim
 c-bsr LoopBack0
#
Return
```

```
[R4]display current-configuration
[V200R001C00SPC500]
#
 sysname R4
#
 board add 0/1 1SA
 board add 0/2 2FE
#
 multicast routing-enable
#
acl number 2000
 rule 5 permit source 225.0.0.0 0.0.0.255
#
acl number 2001
 rule 5 permit source 225.0.1.0 0.0.0.255
#
interface Serial1/0/0
 link-protocol ppp
 ip address 10.0.14.4 255.255.255.0
 pim sm
#
interface GigabitEthernet0/0/1
 ip address 10.0.24.4 255.255.255.0
 pim sm
#
interface LoopBack0
 ip address 10.0.4.4 255.255.255.255
 pim sm
#
ospf 1
 area 0.0.0.0
  network 10.0.14.4 0.0.0.0
  network 10.0.24.4 0.0.0.0
  network 10.0.4.4 0.0.0.0
#
Return
```

```
[R5]display current-configuration
[V200R001C00SPC500]
#
 sysname R5
#
 board add 0/1 1SA
 board add 0/2 2FE
#
 multicast routing-enable
#
acl number 2000
 rule 5 permit source 225.0.0.0 0.0.0.255
#
acl number 2001
 rule 5 permit source 225.0.1.0 0.0.0.255
#
interface Serial1/0/0
 link-protocol ppp
 ip address 10.0.35.5 255.255.255.0
 pim sm
#
interface GigabitEthernet0/0/0
 ip address 10.0.25.5 255.255.255.0
 pim sm
#
interface LoopBack0
 ip address 10.0.5.5 255.255.255.255
 pim sm
#
ospf 1
 area 0.0.0.0
  network 10.0.25.5 0.0.0.0
  network 10.0.35.5 0.0.0.0
  network 10.0.5.5 0.0.0.0
#
pim
 c-rp LoopBack0
#
Return

[S1]display current-configuration
!Software Version V100R006C00SPC800
 sysname S1
```

```
#
 multicast routing-enable
#
acl number 2000
 rule 5 permit source 225.0.0.0 0.0.0.255
#
acl number 2001
 rule 5 permit source 225.0.1.0 0.0.0.255
#
interface Vlanif1
 ip address 10.0.24.1 255.255.255.0
 pim sm
#
interface LoopBack0
 ip address 10.0.11.11 255.255.255.0
 pim sm
 igmp enable
 igmp static-group 225.0.0.1
 igmp static-group 225.0.1.1
#
ospf 1
 area 0.0.0.0
  network 10.0.24.1 0.0.0.0
  network 10.0.11.11 0.0.0.0
#
pim
#
Return
```

在线学习资料支持

您可以在华为企业业务网站获得E-Learning课程、培训教材 、产品资料、软件工具、技术案例等：

1、 E-Learning课程：登录*华为在线学习网站*，进入“*华为培训/在线学习*”栏目

免费E-Learning课：对网站所有用户免费开放

职业认证E-Learning课：通过任何一项职业认证即可学习所有职业认证培训E-Learning课程

渠道赋能E-Learning课：对华为企业业务合作伙伴免费开放

2、培训教材：登录*华为在线学习网站*，进入“*华为培训/面授培训*”， 在具体课程页面即可下载教材。

华为职业认证培训教材、华为产品技术培训教材。无需注册即可下载

3、华为在线公开课(LVC)： http://support.huawei.com/ecommunity/bbs/10154479.html

企业网络、UC&C、安全、存储等诸多领域的职业认证课程，华为讲师公开授课

4、产品资料下载： http://support.huawei.com/enterprise/#tabname=productsupport

5、软件工具下载：http://support.huawei.com/enterprise/#tabname=softwaredownload

更多内容请访问：

- http://learning.huawei.com/cn
- http://support.huawei.com/enterprise/
- http://support.huawei.com/ecommunity/